First Report of the Standing Committee on Upgrade Installation-to-Physics Commissioning (SC-IPC)

October 12, 2004

Table of Contents

1.0 Introduction and Charge to the SC-IPC

The D0 spokespersons have convened a standing committee [1] on upgrade installation to physics commissioning (SC-IPC) which has been charged with the development of an integrated experiment-wide plan for the installation and full commissioning of the upgrade elements being prepared for the RunIIb detector. Chaired by Cecilia Gerber and Richard Smith, the SC-IPC will provide periodic written reports of its findings to the collaboration, and the information will be used to evaluate the optimal approach to, and timing for, the installation of the upgrade elements. The full text of the charge is provided in section 6.0.

Fourteen D0 physicists experienced in RunIIa detector installation and commissioning and/or RunIIb hardware and/or software development comprise the membership of the SC-IPC.

2.0 Working Groups of the SC-IPC

The SC-IPC divided itself into working groups consisting of four or five individuals each and typically each member of the SC-IPC contributed directly to the work of two or more of the working groups. The working groups were defined by the hardware and simulation elements of the RunIIb Upgrade Project [2] and leaders of five were designated. Three members of the SC-IPC (Ela Barberis, Bill Cooper, Bob Hirosky, and Ken Johns) are Level III managers of the Upgrade Project. The name of the leader of each working group is underscored in the following:

2.1 Working Group 1 (L1 Cal Trig)

Ken Johns, Ela Barberis, Volker Buescher, Bob Hirosky

2.2 Working Group 2 (L1CalTrack & L1CTT)

Stefan Grünendahl, Ken Johns, Gordon Watts, Breese Quinn

2.3 Working Group 3 (L2 βeta Upgrade, STT Expansion)

Bob Hirosky, Rick Jesik, Gordon Watts, Taka Yasuda

2.4 Working Group 4 (Trigger Simulation)

Jon Hays, Elizaveta Shabalina, Volker Buescher, Rick Jesik, Eric Kajfasz, Eckhard von Toerne

2.5 Working Group 5 (Layer 0 Silicon)

Breese Quinn, Ela Barberis, Bill Cooper, Eric Kajfasz, Elizaveta Shabalina, Taka Yasuda, Eckhard von Toerne

2.6 Working Group 6 (AFE II)

Working Group Six (AFE II) has not yet been formed. The AFE II evolved from work originally done to prepare the VLPC readout for anticipated 132 ns bunch crossing operation. To replace the SIFT and SVX2 of the AFE I (neither would function at 132 ns) a new trigger pipeline (Trip) chip had been developed and prototyped. When the future of 132 ns operation became unlikely it was realized that the Trip also could be provided with time information when operated at 396 ns bunch crossing to provide z-information for each hit, effectively doubling the number of channels of the CFT (which, for a typical benchmark tracking criterion, would increase the luminosity limit for effective tracking from e.g. 100×10^{30} cm^{-2}s^{-1} to 200×10^{30} cm^{-2}s^{-1}). At the same time, significant limitations in the performance of the existing AFE I at increasing luminosity were observed which would inevitably impair the performance of the VLPC at anticipated RunIIb luminosities. To simultaneously capitalize on the significant investment made during the 132 ns R&D program (the first submission of the Trip chip was found to provide superior operation in every way at 396 ns than the old SIFT/SVX2 combination) and to address the performance limitations of the present AFE I system at high luminosities, D0 has embarked on an aggressive program to prepare a new AFE II, equip it with a revised Trip with timing (Trip-t), and advance the schedule of the work sufficiently so that it might be installed to replace the AFE I perhaps as soon as late 2005. D0 expects a prototype AFE II with Trip-t to become available for testing in late 2004, and at that time the SC-IPC will develop an estimate of the costs in time and effort to install a full AFE II system in the detector.

2.7 Working Group 7 (Online)

No distinct working group will be convened by the SC-IPC for the Online System upgrades. The upgrade tasks for the online system include increasing the Level 3 processing capability, replacing/upgrading host system components, and upgrading as much as 1/3 of the control processors in the experiment control system. This work is expected to be accomplished "adiabatically" largely by the procurement of the most powerful processors as late as possible consistent with running after the general 2005 shutdown. For the host system, clusters of the system are scheduled to be installed and made operational during the 2004 shutdown so that the addition of the remaining processors for data logging can be staged later.

3.0 Executive Summary of the First Report

The SC-IPC finds the collaboration intensively engaged in the development of RunIIb hardware as managed by the Upgrade Project [2]. The collaboration is to a lesser extent engaged in the development of the software needed to operate the hardware. Within the Upgrade Project significant technical changes have been recently made (e.g. to the L1CTT, rescoped March 2004 to include technical improvements to greatly facilitate the commissioning of the L1CTT; or for the installation of a L1CalTrig test area to precommission the full L1CalTrig system using live trigger input from the calorimeter during 2004-2005). Also, a Trigger Steering Committee has been formed to guide the final evolution and implementation of the RunIIb trigger architecture as a whole. [As this report was in final draft the collaboration also appointed a Trigger Coordinator; see section 3.4].

Nevertheless, the totality of installation and commissioning tasks for the several upgrade elements are daunting and a great deal of focused management will be required to ensure that installation and commissioning proceed efficiently so that when the detector resumes physics data-taking there will have been minimal cost in luminosity not logged to tape.

In its work the SC-IPC has assumed that the Upgrade Project will deliver all hardware items "bench-tested" for functionality, and in many cases as will be seen below, operationally tested in the running D0 environment. The SC-IPC understands that the managers of the various upgrade pieces will prudently take advantage of pertinent technical innovations that are encountered along the way to improve the function of any given subelement. Necessarily the SC-IPC has based its evaluation of effort and time required for actual installation, technical commissioning, and physics commissioning on the functionality of each upgrade piece as it is understood at the present time. Furthermore, the Upgrade Project provides a number of key dates corresponding to the completion of each of the upgrade elements. The SC-IPC accepts these dates as "earliest possible dates of completion" and where required, bases its subsequent general planning on them. The SC-IPC will remain aware of unanticipated changes in the Upgrade Project dates and monitor the need for installation planning changes as they occur.

As the SC-IPC has assembled the overall schedule for installation/commissioning it has integrated the manpower needs identified by the Working Groups into the timeline of the schedule. The timelines and effort schedules for the separate upgrade systems are presented individually, and the total effort and timeline presented merely sums these, **assuming as it does that all the upgrades are installed simultaneously in the 2005 shutdown of the Tevatron**. This accountancy might otherwise entail the danger that critical personnel are in fact required by more than one subsystem at the same time. This danger is obviated by the fact that the individual timelines and effort schedules are taken from a single Microsoft© Project schedule for the entire RunIIb Upgrade Installation. This schedule identifies specific personnel where they are known so that overlaps can easily be detected.

In fact the Layer Zero silicon and the trigger upgrades could be installed independently of one another, but it might make little sense for D0 to do this since the Tevatron shutdown required for Layer Zero is quite commensurate with the time the trigger is disabled during installation of the new L1Cal Trig, and the times to return both to physics data-taking are also commensurate.

Table 1 shows the estimates for physicist effort the SC-IPC believes will be required to install and commission the upgrades. The SC-IPC has not identified which of these personnel will automatically be drawn from the Upgrade Project, or which are scheduled to be filled by persons of the appropriate expertise by institutions already committed to the Upgrade. **This means that the SC-IPC has not counted how much of the effort must be newly recruited from the experiment.** Suffice it to say, such effort must come from those able to work for perhaps a dozen weeks or more in certain cases to accomplish the installation and commissioning as planned. The collaboration will realize that it is necessary to ensure that the levels of expertise specified are assigned as required so that return-to-physics will not fall far short of expectations.

In general the SC-IPC finds that members of the Upgrade Project are actively planning to participate, after delivering their particular part of the upgrade hardware, in the installation and commissioning of the upgraded detector to return it physics data-taking as quickly as possible. Such commitments from the Upgrade Project experts will be critical for the success of the installation and commissioning effort.

The methodology employed by the SC-IPC to develop the effort estimates is the customary one. A task is identified and the effort to accomplish it estimated by one or more approaches: scaling from a similar job done earlier, or reckoning the total effort by breaking it down into smaller steps each of which can be estimated, all the while trying to understand the benefits that might accrue from prior experience with similar tasks. This procedure is highly vulnerable to systematic uncertainty since certain of the tasks are understood rather broadly, and many may contain unforeseen and/or inestimable obstacles to success. These uncertainties make it clear that especially the estimates of the time lapses required for the physics commissioning steps are just that. No error bars are assigned to them and they are more likely to be the minima that can be expected rather than the maxima. But the SC-IPC believes the end result of the planning process, however imprecise its individual elements, can guide the experiment meaningfully in planning for the installation and estimating its cost in luminosity forgone that is to be recovered by superior functioning of the detector during Run IIb.

The SC-IPC has exploited the installation timeline to understand what the cost in physics output will be between the time installation begins and the time it concludes at the end of physics commissioning at which point the detector has returned to physics operation at least equal to that it provided prior to the shutdown. Evidently the minimum sum of physics costs is obtained if all the upgrades are installed during a single shutdown. It is found that the installation of the Layer Zero silicon requires the greatest down-time of the Tevatron (8 weeks total, minimum), and the installation of L1CalTrig (the trigger element whose installation makes the greatest impact on data-taking) can be done in the time-shadow of the Layer Zero installation. (L1CalTrig installation does not require shutdown of the Tevatron, but during its installation the luminosity cost to D0 is the same as if it did). Ideally of course, the timing of the Tevatron shutdown should be chosen so both elements are fully ready for installation. Except for perhaps L2Beta, it is apparently optimal to install all the trigger upgrades at one time, again ideally during the same shutdown that is scheduled for the silicon.

After installation the stages of technical commissioning for each of the trigger upgrade pieces, followed by commissioning to physics, extends the time during which the luminosity cost to the physics program of D0 accrues. It is found that while at week eight new L1Cal Triggers can go to the TFW for the first time, another four weeks are required to complete technical commissioning for L1Cal at which time global data-taking is possible. After an additional eight weeks for its physics commissioning the performance of the new system is fully understood and verified and high-luminosity physics data-taking can resume. At week 12 the new L1CTT is technically functional, the new L1CalTrack two weeks after that, and at week 21 both are fully certified for physics data-taking. L2STT is fully integrated into the trigger system and certified for physics data-taking at week 21. Because the RunIIa silicon remains operational, the installation of Layer Zero generates no added luminosity cost even if the return to physics of Layer Zero is delayed. The planning indicates at week 11 data from the Layer Zero silicon can go to the DAQ, and in six additional weeks Layer Zero silicon will be fully certified and its data integrated into the trigger.

From an analysis of the various readiness times, the SC-IPC has estimated the overall physics costs that accrue during the entire period. As the frequency of special runs decreases towards the end of the physics commissioning of the various upgrades, the cost in weekly luminosity is assumed to taper to 50% in the 4th and 3rd final weeks, and to 25% in the 2nd and 1st final weeks before full physics data-taking is achieved. The estimates of the costs in time, effort, and luminosity are summarized in Table 1.

Upgrade Physicist Effort and Luminosity Cost Timeline

Task Name	Wk1	Wk 2	Wk 3	Wk 4	Wk 5	Wk 6	Wk 7	Wk 8	Wk 9	Wk 10	Wk 11	Wk 12	Wk 13	Wk 14	Wk 15	Wk 16	Wk 17	Wk 18	Wk 19	Wk 20	Wk 21	Totals
Tevatron Shutdown																						
Upgrade Activity	Install Hardware							Technical Commissioning				Physics Commission										
Layer Zero Silicon																						
L1 Cal Trig																						
L1 CTT																						
L1 CalTrack																						
L2 Beta																						
L2 STT																						
Trigger List																						
Integrated Luminosity*	0	0	0	0	0	0	0	0	0	0	8	18	30	43	58	75	92	111	131	151	172	
RunIIb Upgrade Cost	0	0	0	0	0	0	0	0	0	0	8	18	30	43	58	75	92	102	112	117	122	
Fermilab Physicists	3	6	2	5	5	5	2	1	1	3	1	1	0	0	0	0	0	0	0	0	0	33
University Physicists	1	5	5	7	5	5	4	4	3	3	3	6	5	5	5	5	2	2	2	2	2	82
Physicists	14	14	10	6	6	12	20	20	16	16	12	16	15	15	18	23	23	25	31	19	19	350
Total	18	25	17	17	16	22	26	25	20	22	16	23	20	20	23	28	25	27	33	21	21	465

* J. Spaulding, BD, 8/04 Design, 10wk shutdown

Table 1. Upgrade Installation Physicists Required, Luminosity Cost Timeline

The anticipated weekly luminosities after a 10-week shutdown are taken from recent schedules provided by the Beams Division. Note that physicists are denoted "University" or "Fermilab" only if the SC-IPC believes that there exists a specific commitment by either Fermilab or a university to contribute to the effort. Where this commitment is unknown to the SC-IPC, "Physicist" is used to designate the required effort.

Not shown in Table 1 are the estimates for the other types of personnel required to accomplish the installation. These estimates for engineers indicate that no more than four persons are required at any one time, for computer scientists two, and for technicians thirteen (peaking in the second week of the shutdown). The numbers of engineers and technicians required is not grossly inconsistent with the present size of the Fermilab D0 operating group.

As will be detailed in the following sections, some "infrastructure" tasks for installation must begin well before the shutdown of the Tevatron. Chief among these are software tasks not delivered by the Upgrade Project and certain tasks that pertain to the precommissioning of trigger elements (especially L1Cal and L1CTT) and those that prepare the detector for the installation of a given upgrade element. It is found that a total of approximately 142 physicist-months ("PM") of physicist effort are required prior to the shutdown. The primary consumers of this effort are the Layer Zero offline software (20 PM), the L1Cal infrastructure and precommissioning (23 PM), L1CTT/L1CalTrack software and precommissioning (21 PM) and trigsim and trigger list development (69 PM).

The SC-IPC is unaware of any study which has quantified the integrated cost to the experiment that the aging of its present silicon will ultimately engender, or the integrated cost accrued by the inability of its trigger to accommodate increased L1 rates as a means of coping with increasing luminosity. Nevertheless the SC-IPC assumes that these costs are real and substantial, and when the benefits of improved function especially of the trigger are added, they more than offset the costs of luminosity shown in Table 1 that are paid because of the installation of the upgrades. Specific examples of the benefits to the D0 physics program afforded by e.g. L1CalTrig are described in section 7.1, or Layer Zero, described in section 7.5.

The contributions made to Table 1 by the individual components of the upgrade are summarized in what follows. Specific recommendations of the SC-IPC for each of the upgrade elements in the following subsections are printed in italics.

3.1 Working Group 1: L1CalTrig

The Upgrade Project manages the L1CalTrig upgrade to a hardware date-of-completion of July 1, 2005. There are four distinct stages for returning to physics data-taking with the L1CalTrig upgrade: precommissioning, installation, technical commissioning, and physics commissioning. Of these, the precommissioning stage is by far the most important since most of the hardware integration, software development, and trigger studies must be accomplished during this period. Assuming all precommissioning goals are met, the SC-IPC believes L1CalTrig will be ready for global data-taking four weeks after installation and for physics data-taking four - eight weeks after that. The SC-IPC finds that significant effort will be required to achieve the precommissioning goals.

Since late CY 2003 the L1CalTrig group has operated a test stand adjacent to the D0 Moveable Counting House (MCH1) where trigger signals split off from the CAL BLS have been available for testing of L1CalTrig components. Between the hardware date of completion and the beginning of installation, the group will gain valuable experience operating the entire L1CalTrig system in this precommissioning environment, albeit with a small number (32) of CAL BLS inputs.

The precommissioning work outside MCH1 is critically important to the success of L1CalTrig. First, all of the trigger hardware must be integrated and readout into the experiment. Second, all of the control, online, and offline software for operating the system must be developed during this period. Third, a large amount of trigger certification must be performed. This includes the correct simulation of the trigger tower (TT) response as well as optimization and certification of trigger algorithms. Certification demonstrates that the trigger simulator output correctly models the hardware to better than 1-2%. Fourth, trigger studies must be carried out in order to produce a post-shutdown trigger list for physics data-taking.

There is the danger that among the variety of associated software and trigger certification tasks some will be overlooked while the hardware is completed and integrated into the experiment. The time penalties for not achieving the full goals of the precommissioning work are substantial and will be applied the technical and physics commissioning periods, thus delaying the return to physics data-taking. During the latter parts of the precommissioning period regularly scheduled shifts sustained in the test area will ensure that precommissioning is not abbreviated.

When the actual time for installation is chosen, the group plans to decable and remove the RunIIa trigger system in MCH1, roll in the racks of the new system from the testing area, connect the crates via patch panels to the existing BLS cables, and begin the technical and physics commissioning of the new system.

Following installation another intense period of activity must occur to insure a rapid return to physics data-taking. A brief technical commissioning period is needed to re-establish operation of the full L1CalTrig system and to check all inputs. This will evolve into the physics commissioning period where collider data will be used to tune the ADF digital coefficients, perform final trigger rate studies, and confirm the earlier trigger certification. The teams that worked on completion of the hardware and its precommissioning in the test area must be augmented with other experts to perform the many physics commissioning tasks.

Table 2 presents the timeline and effort summary required to take the L1CalTrig from the beginning of installation to the point that it is fully operational in high luminosity physics data-taking. Seven weeks are required for installation, four weeks for technical commissioning, and up to eight weeks for physics commissioning. The SC-IPC believes that global data-taking can resume after the technical commissioning period and that physics data-taking can resume when physics commissioning ends. To repeat, any work or studies not faithfully completed during precommissioning will add additional time to these estimates.

Note that the TFW is off at the beginning of the period for only ~ two days, and at week eight L1CalTrig is producing output to the TFW (certified output by week 12) for debugging other components of the upgrade.

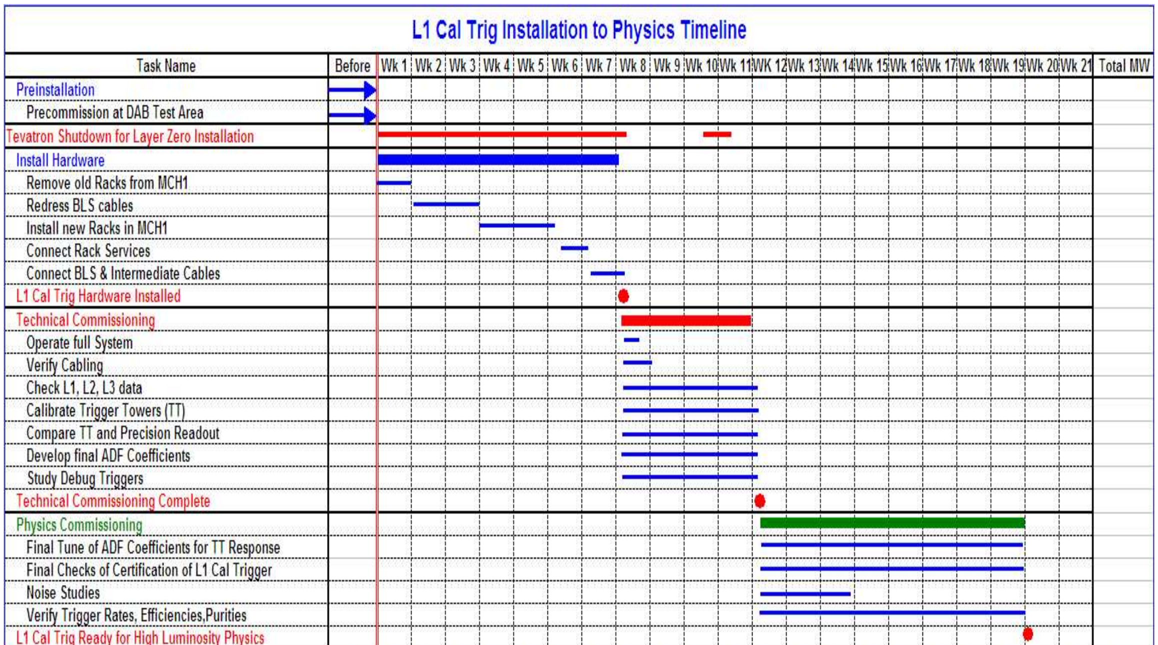

Task Name	Before	Wk 1	Wk 2	Wk 3	Wk 4	Wk 5	Wk 6	Wk 7	Wk 8	Wk 9	Wk 10	Wk 11	WK 12	Wk 13	Wk 14	Wk 15	Wk 16	Wk 17	Wk 18	Wk 19	Wk 20	Wk 21	Total MW
Fermilab Physicists	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Fermilab Technicians	0	7	8	3	3	4	6	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	35
University Physicists	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2
Physicists	4	4	4	4	0	0	6	6	6	6	6	3	5	7	7	10	10	10	10	10	0	0	114
Total Physicists	4	4	4	4	0	0	7	6	6	6	6	3	5	7	7	10	10	10	10	10	0	0	116

Table 2. L1CalTrig Installation to Physics Timeline & Effort Summary

Because L1CalTrig can provide triggers as soon as it is technically commissioned, other portions of the trigger upgrade and the Layer Zero silicon can be commissioned for physics as soon as the physics commissioning of L1CalTrig begins in week 12. No useful physics data will be taken during the technical commissioning of L1CalTrig, and physics data collected during the physics commissioning period may not be useful for all physics analyses. It might be possible for example to utilize data collected with b-physics triggers during this period, or longer if L1CalTrig commissioning is protracted, but the SC-IPC does not attempt to quantify this possibility.

Note Table 2 begins prior to the installation shutdown with a column labeled "Before" which emphasizes the effort required to conduct the large list of precommissioning tasks. These include hardware precommissioning, software development, and trigger certification not provided by the Upgrade Project. The schedule indicates this sums to 90 physicist-weeks of effort, to be added to the total column of Table 2.

Note also that physicists are denoted "University" or "Fermilab" only if there exists a specific commitment by either Fermilab or a university to contribute to the effort. When this commitment is unknown to the SC-IPC, "physicist" is used to designate the required effort.

The SC-IPC recommends that the collaboration identify a position which will be responsible for the oversight and coordination of the L1CalTrig precommissioning effort. This position may well require two individuals working together, one to coordinate the precommissioning hardware infrastructure and another to coordinate the online and offline software development and data *analysis that underpins the entire enterprise. The SC-IPC believes there is significant danger to the upgrade schedule if no one is made responsible for these two areas. The SC-IPC further believes that the magnitude of the tasks they face requires that the position(s) be given prominence on the appropriate organization chart of the experiment. The position(s) will profit from this prominence by the visibility thereby given to the precommissioning and installation efforts. This visibility will enable the manager(s) to articulate to the collaboration the specific recruiting needs for the long term groups and short term members that are required to complete the installation of L1CalTrig successfully. This recruitment can be aided by the recognition that certain of the software tasks will satisfy the service requirements of postdocs and graduate students.*

The SC-IPC also recommends that manpower for regularly scheduled shifts in the test area during the precommissioning period be managed from the existing D0 shift pool (i.e. those taking the shifts receive shift service credit for their institutions).

3.2 Working Group 2: L1CalTrack & L1CTT

The Upgrade Project manages the L1CalTrack trigger to a hardware date-of-completion of May 16, 2005. Well in advance of this date, i.e. during the Summer 2004 shutdown, the (empty) L1CalTrack crate and power supply will be installed in MCH1. After the run resumes in the fall of 2004, this substantial infrastructure will facilitate the technical precommissioning of L1CalTrack prior to the 2005 shutdown. The L1CalTrack crate will be initially populated with spare L1MU trigger and crate manager cards. Spare L1Mu inputs can be used to generate BOT triggers to send to the TFW, and observing a stable 47712 Hz BOT rate will demonstrate overall system functionality. Between or during stores the L1CalTrack crate can used to check outputs to L1 (i.e. the TFW) and L3. This provides a working trigger system albeit one with only a few inputs. The operation of this system allows development and testing of online and offline software needed for L1CalTrack. As production L1CalTrack trigger and crate manager cards become available they will be swapped for the L1MU cards. At this point, further commissioning must await the arrival of the new L1CalTrig and L1CTT inputs. The development of the necessary software for the operation of the new L1CalTrack (downloads, databases, operator gui's, etc.) must occur as the new system is precommissioned so that when the shutdown begins the on-time physics commissioning of the new system does not incur large delay. See chapter 4.2.1 for details.

The Upgrade Project manages the L1CTT upgrade to a hardware date-of-completion of May 23, 2005. To facilitate precommissioning the new system the L1CTT upgrade also includes the advance installation of an LVDS splitter system to enable the in-situ operation of the new DFEB's as they emerge from production. Thus during the summer 2004 shutdown, the LVDS splitter crate, a temporary (prototype) DFEA crate plus one additional CTOC and CTTT board, and the new 48V power supply, will be installed on the platform (which with a few VRB channels in crate 0x13 provides a complete secondary CTT readout chain). As they become available the new crate controller and prototype DFEB's can be loaded in the temporary crate and tested with real data in parallel to the RunIIa CTT. At the beginning of the 2005 shutdown, the RunIIa CTT crates will be removed from the platform, the two new crates installed and cabled, the 48V power run to the rack, and technical commissioning completed as the beam returns following the installation of the Layer Zero silicon.

The schedule for the installation of the LVDS splitter system (in the 2004 shutdown) is critical for the precommissioning for L1CTT. The development of the necessary software for the operation of the new L1CTT (downloads, databases, operator gui's, etc.) must occur as the new systems are precommissioned in parallel with the operation of the existing L1CTT, and the penalties to the on-time commissioning of the new systems once they are installed is large if this effort is not comprehensive. See chapter 4.2.1 for details.

Table 3 presents the timeline and effort summary required to take the L1CTT and L1CalTrack from the beginning of installation to the point where they are fully operational in high luminosity physics data-taking. It is seen that from the start of the shutdown (by assumption, for the installation of the Layer 0 silicon), one week elapses before the detector platform is accessible (while the detector is being opened up) and hardware installation for the L1CTT begins. The L1CalTrack hardware is by assumption already installed in MCH1. Because of the precommissioning given to both systems, technical commissioning consists largely of debugging inputs and verifying outputs of the two systems. The later stages of technical commissioning for L1CalTrack must wait until L1CTT and L1CalTrig become available. Then when technical commissioning concludes, physics commissioning occurs for L1CTT and L1CalTrack.

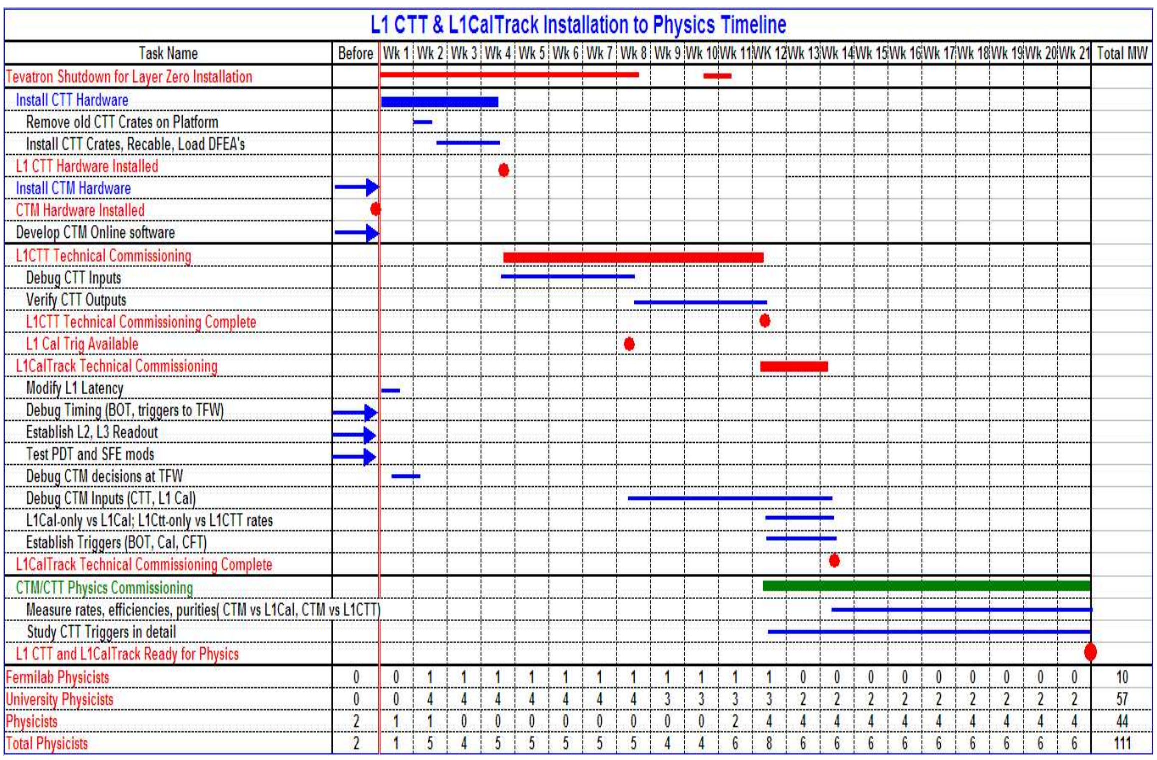

Table 3. L1CTT & L1CalTrack Installation to Physics Timeline & Effort Summary

As seen in Table 3, approximately 9-10 weeks elapse after the L1CTT is technically commissioned until the L1CTT and L1CalTrack are contributing fully to physics. Technical commissioning of L1CalTrack properly takes place after L1CTT and L1CalTrig are operational (but L1CalTrig could delay 3-4 weeks before the physics commissioning of L1CalTrack became the pacing item. While it could be said that L1CalTrack is not crucial for return to physics data-taking, L1CTT is, so it is appropriate to conclude that except for the last ~ 4 weeks of the period all delivered luminosity is lost to the experiment. During these last four weeks it can be imagined that special runs for the physics commissioning might consume only 50% of the time in the first two weeks, and only 25% of the time in the last two. If this proves realistic, the luminosity cost to the experiment will be reduced.

Note that table Table 3 begins with a pre-installation column labeled “Before” which emphasizes the fact that the L1CalTrack hardware and online software be fully installed and debugged before the general installation shutdown. Also, during this period, substantial parts of the technical commissioning of L1CalTrack must be achieved. Table 3 indicates that two physicists beyond those provided by the Upgrade Project are required to assure that the precommissioning of L1CalTrack and L1CTT proceed as required. The details of this effort are described in section 4.2. The schedule indicates this sums to 83 physicist-weeks of effort, to be added to the total column of Table 3.

Note also that physicists are denoted “University” or “Fermilab” only if the SCD-IPC is aware that there exists a specific commitment by either Fermilab or a university to contribute to the effort. When this commitment is unknown to the SC-IPC, “physicist” is used to designate the required effort.

3.3 Working Group 3: L2 βeta Upgrade & STT

The Upgrade Project manages the L2 βeta upgrade to a hardware date-of-completion of March 24, 2005. Well in advance of the shutdown to install the other trigger upgrades, new L2 βeta processor cards will be installed in the L2 trigger framework on an “adiabatic” basis as they become available. A scheme for bit expansion at L2 is expected to have been completed and installed in the trigger system prior to the 2005 shutdown as well. Once the other trigger components become operational, software changes to accommodate the data format change of the L1CalTrig upgrade must be commissioned, and data verification made on data from all other upgraded processor systems.

The Upgrade Project manages the L2 STT upgrade to a hardware date-of-completion of April 1, 2005. Well in advance of the shutdown to install the other upgrades, additional splitters and optical fibers will be installed in the MCH for the new STT. At the start of the installation shutdown, the additional STT hardware will be installed. Data flow from the L0 silicon will be verified as soon as it is available and properly synchronized with SCL and L1CTT inputs. Also as soon as the new L1CTT becomes available, L2CTT data flow will be verified.

Table 4 presents the timeline and effort summary required to take the L2 βeta and STT from the beginning of installation to where they are fully operational in high luminosity physics data-taking. It is seen that from the start of the installation shutdown all new hardware can be installed in one week. Because the STT infrastructure (splitters and optical fibers) have by then already been installed in the MCH, the technical commissioning of the new systems will begin as soon as the schedules for the installation of the L1CTT and Layer Zero silicon allow. Since the STT upgrades are essentially augmentations of existing systems, and the tracking algorithms will have been developed for trigsim, physics commissioning is expected to be straight-forward.

L2 Beta & STT Installation to Physics Timeline																						
Task Name	Wk 1	Wk 2	Wk 3	Wk 4	Wk 5	Wk 6	Wk 7	Wk 8	Wk 9	Wk 10	Wk 11	Wk 12	Wk 13	Wk 14	Wk 15	Wk 16	Wk 17	Wk 18	Wk 19	WK 20	Wk 21	Total MW
Tevatron Shutdown for Layer Zero Installation																						
Install Hardware																						
Install L2 Beta processors in L2 Trigger																						
Install STT Modules																						
Install VTM Modules																						
Install TFC modules																						
L2 Beta and STT Hardware Installed																						
Technical Commissioning																						
L2Beta Technical Commissioning																						
Layer Zero Functional																						
Layer Zero Technically Commissioned																						
L1CTT Technically Commissioned																						
STT Technical Commissioning																						
Technical Commissioning Complete																						
Physics Commissioning																						
Verify New Trigger Bit Expansion, etc.																						
Verify STT Track Fitting Code																						
L2 Beta & STT Ready for High Luminosity Physics																						
Fermilab Physicists	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	3
University Physicists	1	2	2	2	0	0	0	0	0	0	0	3	3	3	3	3	0	0	0	0	0	21
Physicists	3	3	0	0	0	0	0	0	0	0	0	0	0	0	0	4	4	4	4	4	4	30
Total Physicists	5	6	2	2	0	0	0	0	0	0	0	3	3	3	3	7	4	4	4	4	4	54

Table 4. L2 βeta and STT Installation to Physics Timeline & Effort Summary

As is seen in Table 4, the beginning of technical commissioning of STT is conservative. Likely the STT experts will begin looking at outputs from Layer Zero and L1CTT before the technical commissioning of these is fully complete. The beginning of physics commissioning of STT is also conservative and likely can start two weeks or so earlier for the same reason. As with L1CalTrig, L1CTT and L1CalTrack, it is likely that the return to physics will be gradual, i.e. during the last four weeks of the table, special runs will not consume more than 50% of the time the first two weeks, and 25% of the time in the last two weeks. If this proves realistic, then the total luminosity cost to the experiment will be reduced.

Note that physicists are denoted "University" or "Fermilab" only if there exists a specific commitment by either Fermilab or a university to contribute to the effort. When this commitment is unknown to the SC-IPC, "physicist" is not designated as belonging to either.

There is a minor amount of effort required (7 physicist-weeks) prior to the general shutdown of the Tevatron that must be added to the totals of Table 4.

3.4 Working Group 4: Trigsim

At the resumption of physics following the installation of the upgrades, it is necessary to have a RunIIb Trigger List in which there is high confidence that all desired triggers from the old and new parts of the detector are properly collected, and that the prescale sets are at least reasonably efficient at the outset of running.

A trigger simulation expert tool ("*trigsim*") with ancillaries ("*trigsimcert*" and "*ratetool*") developed for RunIIa have proven invaluable in constructing useful trigger lists and understanding the physics analysis trigger studies using Monte Carlo and detector raw data. These expert tools must be upgraded to accommodate the changes in e.g. L1CalTrig, L1CalTrack, the new Layer Zero silicon, and L2STT.

The Upgrade Project makes crucial contribution to the development of the RunIIb trigger list in its efforts to integrate the new L1Cal Trig, L1CalTrack, and L1CTT into the trigger. The trigger group of the Upgrade is optimizing algorithms for RunIIb, preparing individual trigger upgrade simulator components (*trigsim*), and working to upgrade *ratetool* so algorithm choices can be made and the L1 terms of the trigger list can be developed. At the July 2004 Director's Review of the Upgrade, the group was encouraged to make a formal presentation to the Director in early 2005 on the trigger simulator and validation efforts. The Upgrade group plans to complete the *trigsim* modifications in "early" 2005. (The Upgrade schedule presently indicates L1CalTrig and L1CTT will be integrated into trigsim by the end of 2004, and L1CalTrack by March 2005). L1CalTrack might be finished earlier. Progress with the L2STT simulator must wait until Layer Zero is integrated into the software, and this latter activity has not yet started.

As detailed in section 4.4 the SC-IPC developed a schedule for having in place a RunIIb trigger list not later than the time the trigger upgrades were installed and physics commissioning was underway, so the return to physics-quality data-taking was not delayed. Working back in time from this optimum milestone and incorporating the best experiences from the development of earlier trigger lists, it was found that a straw man trigger list no later than February 2005 was necessary, as well as full integration of the upgrade elements into *trigsim* with *trigsimcert* and *ratetool* fully operational so that the trigger list could be certified no later than March 2005. Then the full L1/L2 trigger menu would be finalized in the following 4—5 months, L3 another month after that. Then after 6 weeks of data base entry and checking, the first physics data could be collected with the trigger list.

The SC-IPC believes that this schedule is aggressive and the collaboration must dedicate the effort detailed in section 4.4 starting immediately, plus address the likely need to accelerate the integration of the upgrades into *trigsim* so confirmation of the strawman trigger list is not delayed and all the necessary trigger rate studies are made. It is particularly important that those who join in the effort expect to interface seamlessly with those already working in the Upgrade project so not a motion is lost or task duplicated. Failure to achieve the proposed schedule will cause the return to physics to be delayed and the cost in luminosity to increase accordingly.

The SC-IPC endorses the recent action of the Spokespersons to identify a person who will be responsible for the integration of the RunIIb Trigger effort in the overall D0 Trigger strategy. (The creation of a position of this type was recommended in drafts of this report that predated the action of the Spokespersons).

Furthermore the SC-IPC recommends that this person be assisted by a group of individuals composed of certain of the present users of the various Trigsim packages (e.g. with representatives from each of the physics groups in addition to trigger and simulation specialists), and together with its manager the group be charged with developing the RunIIa trigsim into a robust, user-friendly package intimately connected to Monte Carlo and RECO, with a toolkit of resources necessary to verify its operation and validity, that this group assemble the RunIIb trigger list, that it conduct the trigger rate studies necessary to validate the trigger list, and that this group continue its longevity past the beginning of the Run IIb era at which time Trigsim is a full partner in the expression of the Run IIb trigger.

Table 5 presents the timeline and effort summary required to develop trigsim and prepare a trigger list and prescale set for use after the installation and commissioning of the upgraded hardware is complete. By assumption, the first week of November 2005 corresponds to week 20 of the installation shutdown.

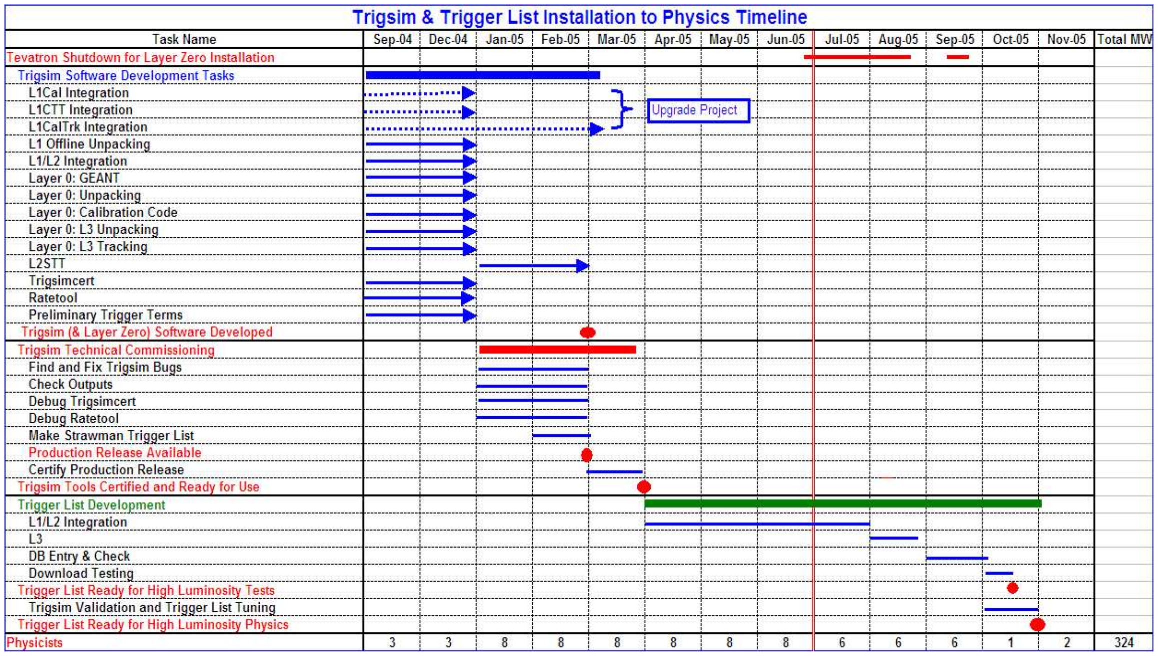

Table 5. Trigsim and RunIIb Trigger List Timeline & Effort Summary

As is seen in Table 5, the effort required for the development and commissioning of the first RunIIb trigger list is substantial, but critical for the success of the upgrade. Provided the work ramps up to the specified levels and is sustained as required from now until the upgrade shutdown is over, it is within the grasp of the collaboration to begin to take physics quality data as soon as the hardware is ready. With the shutdown of the Tevatron as shown, the beginning of November corresponds to week 20 of the shutdown. The analysis shows there is not a moment to lose, especially with respect to the Layer Zero software.

Note that the table is scaled by calendar quarters in the first two columns and monthly thereafter. The effort total shown includes all the effort to the left of the red bars at the beginning of July 2005. The schedule indicates the pre-shutdown effort sums to 146 physicist-weeks of effort.

Note also that "physicists" does not carry the descriptive "University" or "Fermilab" since the SC-IPC hasn't verified formal commitment by either Fermilab or any university to contribute to specific parts of the trigsim effort.

3.5 Working Group 5: Layer Zero Silicon

The Upgrade Project manages to a hardware date-of-completion of July 21, 2005 for the Layer Zero silicon. Because it takes time to open the detector and remove the H-disks before the Layer Zero silicon can be installed, the Tevatron shutdown should begin no sooner than approximately 3 weeks before the silicon is completed at SiDet.

Certain pre-installation tasks (modifying existing online and offline software, and installing certain hardware infrastructure at the detector) must be completed before Layer Zero installation begins. Once Layer Zero is installed, connected, and cooling reestablished, its technical commissioning can begin. Operation of all channels is tested, as is the new online software prepared for Layer Zero. Collisions can resume during this period, but as technical commissioning concludes, collisions must be suspended briefly and the detector closed up and surveyed. At this point collisions resume and physics commissioning begins. The details of these activities are described in section 4.5.1.

Table 6 presents the timeline and effort summary for the installation and commissioning of the silicon.

WG 5: Layer Zero Silicon Installation to Physics Timeline

Task Name	Before	Wk 1	Wk 2	Wk 3	Wk 4	Wk 5	Wk 6	Wk 7	Wk 8	Wk 9	Wk 10	Wk 11	WK 12	Wk 13	Wk 14	Wk 15	Wk 16	Wk 17	WK 18	Wk 19	Wk 20	Wk 21	Total MW
Preinstallation																							
Hardware Infrastructure																							
Online Software																							
Offline Software																							
Tevatron Shutdown for Layer Zero Installation																							
Install Hardware																							
Open Detector, Install Scaffolds, Disconnect Tev Beampipes																							
Warm silicon, Remove EC beampipes, H Disks																							
Install Layer Zero Mounts																							
L0 Ready at SiDet																							
Install Layer Zero, Connect Be Pipe																							
Cable Silicon, Install H Disks, Install Cooling																							
Reconect Tev Beampipes & Leak Check, Cooldown Si																							
L0 Silicon Installed																							
Technical Commissioning																							
Demonstrate Full Operability of all channels																							
Verify Functionality of all SMT Software																							
Close Detector																							
Technical Commissioning Complete																							
Physics Commissioning																							
Timing Studies																							
Noise Studies																							
Alignment																							
Clustering Studies																							
Tracking Studies																							
Physics Objects (Neutral Vee's, etc)																							
L0 Silicon Ready for High Luminosity Physics																							
Fermilab Physicists	1	2	4	1	3	4	4	1	0	0	2	0	0	0	0	0	0	0	0	0	0	0	20
University Physicists	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2
Physicists	3	0	0	0	0	0	0	8	8	4	4	6	6	3	3	3	3	3	3	9	9	9	81
Total Phyicists	4	2	4	1	4	5	4	9	8	4	6	6	6	3	3	3	3	3	3	9	9	9	103
Technicians	0	4	4	6	5	8	6	5	0	0	4	4	0	0	0	0	0	0	0	0	0	0	46

Table 6. Layer Zero Silicon Installation to Physics Timeline & Effort Summary

As is seen in Table 6, the timeline required for the installation of the Layer Zero silicon requires a shutdown of the Tevatron of approximately 10 weeks total, minimum. Ideally the Tevatron could resume in week 9 with closeup of the detector delayed until week 11 to facilitate access to the silicon during technical commissioning. In principle the improvement provided by Layer Zero to tracking and vertex resolution is not necessary for physics to resume, so the SC-IPC does not charge the layer zero installation with any cost in luminosity after week 11 in the schedule unless its installation is delayed beyond the given shutdown of the Tevatron. If for example the Tevatron shutdown is fixed at 10 weeks to accommodate BD or CDF schedules (and D0 closes up promptly to accommodate this resumption), then layer zero installation could cause no cost in luminosity.

Following technical commissioning, physics commissioning with collider data occurs. After noise studies begin, timing-in is done, and when noise studies are complete offline alignment studies begin using data from Layer Zero. After about six weeks of physics commissioning when alignment is well understood and overall data quality is confirmed, the collaboration is assured that all subsequent data is of physics quality. At this time the expected benefits to tracking and vertex resolution afforded by Layer Zero begin to accrue to D0. Further physics analysis continues for some weeks to conduct clustering and tracking studies, and to confirm the presence of physics objects (neutral Vees, etc.) in the data.

Note that Table 6 begins with a column labeled “Before”, which is intended to denote an extended period of time ending just prior to installation, during which three of the approximately four physicists tabulated are engaged in online/L3/offline software development for Layer Zero. Not included in this effort is that required for trigsim work, described in section 4.4, and totaled therein. The software activity must begin immediately and continue until all the software packages itemized in section 4.5.1 and 4.4 are ready for use The majority of the physicist effort in Table 6 is identified only as “physicist”, i.e. not necessarily “Fermilab” or “University”. This implies that the SC-IPC believes those who must contribute have not yet been identified. The integrated effort in the “Before” column is 106 physicist-weeks, to be added to the Total column in Table 6.

The SC-IPC is unaware that anyone is presently identified to coordinate and manage the development of the totality of the Layer Zero software. Given that development of the offline and simulation software must begin immediately the SC-IPC recommends that the collaboration identify a position which will be responsible for the oversight and coordination of the Layer Zero software effort. The SC-IPC believes there is significant danger to the upgrade schedule if no one is made responsible for the Layer Zero software. The SC-IPC further believes that the magnitude of the task requires that the position be given prominence on the appropriate organization chart of the experiment. The position will profit from this prominence by the visibility thereby given to the Layer Zero software effort. This visibility will enable the person who fills the position to articulate to the collaboration the specific recruiting needs for the physicists that are required to complete the software for Layer Zero successfully. This recruitment can be aided by the recognition that certain of the software tasks will satisfy the service requirements of postdocs and graduate students.

3.6 Online

As described in section 2.7 the upgrades of the Online systems consist largely of replacing aging processors with faster more modern components, and incorporating new operating system software where required, in an “adiabatic” fashion with no measurable upset to the experiment.

Perhaps the most invasive of the online tasks is that of the upgrade for the host system. The host nodes perform the functions of data collector, data logger, data distributor, and data shipping to Feynman Computer Center. The present host computers running Compaq Tru64 Unix are aging and face increasing costs of 24x7 maintenance required to ensure >99% ontime availability. They will be replaced with Linux servers clustered to supply the services needed by the host with a Shared Fiber Channel storage and failover software package to provide high flexibility and availability. To reduce the risk of this upgrade, four clusters will be constructed in the fall 2004 shutdown and the entire host (DAQ, Oracle, NFS, etc.) will be moved to the new system. Prior to the general upgrade shutdown in 2005, the best processors available will be purchased to complete the system.

The upgrade of the control system is expected to be truly noninvasive -- ~20% of the processors in the control room are to be replaced each year after 2004 with modern processors and displays.

The upgrade of L3 is also expected to be noninvansive – the L3 system must handle an increased event size and a doubling of the L3 accept rate (50 Hz increasing to 100 Hz) with increasingly sophisticated reconstruction filters for L3 rejection. To achieve these demands, 48 dual 1-GHz nodes of the old system will be replaced with 96 new dual 2.1-GHz nodes which will increase the existing 326 GHz capability of the present system to 634 GHz (e.g. 1 kHz L3 accept rate at 500ms per event requires 500 GHz of CPUs). Procurement of the nodes will be started in time to ensure their availability for the 2005 shutdown. Disruption during the shutdown will be minimized by the fact that the remaining 66 nodes will remain operational during the shutdown.

The Online upgrade is entirely the purview of the Upgrade project and is expected to make no demands for installation effort during the general 2005 shutdown beyond that already scheduled. It is managed and will be accomplished by the manager of the present online system (Stu Fuess) with appropriate support from the Computing Division.

The SC-IPC believes that the installation of the online will generate no cost in luminosity due to delaying the resumption of physics data-taking after the 2005 shutdown.

4.0 Detailed Reports of the Working Groups

In fulfillment of the charge (see section 6.0) to the SC-IPC, each working group has reviewed the technical aspects of the upgrade element it addresses, taking into account the planning that must be done for the physics commissioning of that element and the tasks that must be executed to fulfill the planning, in many critical areas such as:

- Installation and technical commissioning of the hardware;
- Calibration databases, both online and offline;
- Updates to data unpacking & formatting;
- Development of Level 3 filtering algorithms;
- Updates to clustering and track finding, both online and offline;
- Development of Monte Carlo and TRIGSIM.

Each working group has provided an itemization of the effort and time required to accomplish each task it identifies. The interdependencies of the tasks have been outlined and the skill levels of those best able to carry out the work in the time allotted have been defined, often by quoting the name of the person whose expertise was developed in Run IIa on a similar task.

The SC-IPC as a whole has reviewed the separate itemizations of each of the working groups to discern the interrelationships that constrain the separate timelines for the commissioning of the elements delivered by Upgrade Project. From the separate timelines and effort lists from each working group, a sum can be made that defines the totality of effort that will be required to complete the commissioning of the entire upgrade in the time desired.

This same individual timelines have been analyzed by the SC-IPC in an attempt to quantify what the loss of physics is during the commissioning period for the same elements. This latter summarization is complex and will change as progress in understanding the totality of the upgrade evolves within the SC-IPC.

The separate elements of the Upgrade Project are described in detail in available documents for the Layer Zero silicon [3], the L1CalorimeterTrig [4], the L1CTT [5], and the L2 Beta and L2 STT upgrades [6,7]. In section 7.0 brief summaries of the purposes of the upgrade elements are provided to establish a setting for the work of the SC-IPC.

4.1 Working Group 1: L1CalTrig

4.1.1 *Effort & Schedule of Installation and Commissioning*

The four steps to returning to physics data-taking with the L1CalTrig system are precommissioning, installation, technical commissioning, and physics commissioning.

The effort required to precommission the new system at the test site on the sidewalk outside MCH1, and to write all the associated online and offline software, is large. At the outset, after the ADF becomes available, regular two-per day shifts by physicists dedicated to the task should be established in the test area to complete digital filter and trigger certification studies. Only by a constant effort devoted to this precommissioning is steady-state pressure generated to find and solve problems. During this period L1Cal simulation leading to AND/OR terms, D0sim of the ADF response, and development of the offline software can proceed. The effort required for this operation is counted in the totals developed by the SC-IPC.

The development of the L1CalTrig hardware and the software to operate it is managed by the Upgrade Project, and these activities are not the purview of the SC-IPC. In consultation with the L1CalTrig Upgrade Project a detailed precommissioning plan has been developed that includes the enumeration of the totality of the software/analysis tasks that must be completed before installation. The manpower estimates made for these tasks substantially exceed the total present size of the groups building and testing the hardware. Among the software/analysis tasks anticipated for precommissioning are:

a) Demonstrate data acquisition using production ADF, TAB, and GAB boards (SDAQ)
b) Send data to L1, L2, and L3 trigger systems during physics data-taking
c) Develop and implement downloading and system control software
d) Develop and implement online Monitoring Tools (EPICS)
e) Compare TT (trigger tower) and precision readout
f) Compare TT response in data and Monte Carlo
g) Perform digital filter studies
h) Develop calibration procedures
i) Develop online examines
j) Develop offline software (unpacking, DST, TMB, Root-tuples, analysis)
k) Complete tsim_l1cal2b software
l) Continue trigger algorithm optimization
m) Continue L1Cal trigger rate and efficiency estimate studies
n) Continue development of L1Cal physics triggers for data-taking
o) Develop L2 and L3 trigger algorithms

The effort for these tasks must be provided simultaneously while operations at the test site and hardware production and testing is underway. The resources for items a—d in the above list are overseen by the Upgrade Project. The SC-IPC has tabulated the effort and timelines only for the tasks e) and beyond in its analysis of the precommissioning effort required for RunIIb.

Installation of the L1CalTrig consists of removing the existing L1Cal trigger racks from MCH1 and moving the fully populated racks of L1Cal trigger electronics from the test site to MCH1. In somewhat more detail, the BLS cables must be disconnected and lowered into the floor of MCH1, the RunIIa L1CalTrig racks must be removed, the new RunIIb patch panel racks and L1Cal trigger racks moved into MCH1 from the test area on the sidewalk, the BLS cables redressed to the patch panels racks, the ADF cables connected, all racks connected to the MCH rack infrastructure system (cooling, monitoring, safety), and the new L1Cal trigger electronics powered and prepared for commissioning.

After physical installation is complete, a four week long technical commissioning period will follow. Goals of the technical commissioning include reliable operation of the entire system, verifying L1, L2, and L3 data, and checking all inputs. Physics commissioning tasks include tuning the trigger tower (TT) response of all towers and performinag final trigger studies based on collider data. Given that time to install the L1Cal trigger is expected to consume most of shutdown period, final commissioning must occur as quickly as possible in order to regain high efficiency data-taking after the Tevatron resumes operation. This underscores the great importance of a successful precommissioning effort. The danger of an extended physics commissioning period is high. The estimates made by SC-IPC of the timeline and effort for commissioning assume a successful precommissioning effort on the MCH1 sidewalk, and any hardware or software test or debugging not accomplished during the precommissioning effort will significantly delay physics data-taking.

An initial list of technical commissioning tasks includes:

a) Operation of the full L1CalTrig system
b) Verify cabling
c) Check L1, L2, L3 data
d) Calibrate the TT (trigger towers)
e) Compare TT and precision readout
f) Develop final ADF coefficients
g) Study debug triggers

An initial list of physics commissioning tasks includes:

a) Final tuning of ADF coefficients to give correct TT response
b) Final checks of the certification of L1Cal trigger
c) Perform noise studies
d) Verify physics trigger rates, efficiencies, and purities with collider data

A resource loaded schedule for the precommissioning, installation, and technical and physics commissioning of the L1Cal trigger exists and it forms the basis of the SC-IPC analysis presented herein. The results of this analysis is tabulated in the Executive Summary, section 3.1.

4.2 Working Group 2: L1CalTrack and CTT

4.2.1 *Effort & Schedule of Installation and Commissioning*

Installation of the L1CalTrack hardware is abbreviated since by assumption the components of the upgrade have been "adiabatically" installed in D0 before the general upgrade shutdown. Precommissioning involves operating first with L1Mu boards then with L1CalTrack boards as they become available, to demonstrate proper operation of the hardware and all software required to operate it. Technical commissioning begins when the L1CTT and L1CalTrig are available to provide inputs. As inputs become available, resistance and parity measurements are used to verify them. Other technical commissioning tasks include comparison L1Cal-only and L1CTT-only triggers with those generated by the L1CalTrack and L1CTT trigger systems. The most important technical commissioning task is certification which compares the L1CalTrack decisions in hardware with those produced by the trigger simulator using L1CalTrig and L1CTT triggers as simulator input.

During certification of the L1CalTrack trigger, physics commissioning can begin. This consists initially of special runs to measure rates, efficiencies, and purities. Efficiency results can be compared to those of the simulator to bolster confidence in the system. Once L1CalTrack certification and physics commissioning is complete the L1CalTrack trigger should be ready for use in physics data-taking.

Installation of the L1CTT consist of decabling and removing the old crates on the platform, installing the new ones, cabling the new system and then beginning technical commissioning. Specifically, when the operating point for the new system is established and new DFEB crates are ready, the DFEA crates are removed, the DFEB crates are installed, and basic operation for all cards on the platform is established including full CTT CFT axial + CPS system operation. Technical commissioning includes debugging the inputs and verifying the outputs. Eight weeks are scheduled for these activities beginning during week four of the shutdown. Physics commissioning (making examine and offline checks of trigger performance) then proceeds in step with that for L1CalTrak and also concludes at the end of week 21 after the beginning of the shutdown.

The development of the L1CTT and L1CalTrack hardware is managed by the Upgrade Project, and these activities are not the purview of the SC-IPC. However, substantial software work is required to enable these upgrades to be commissioned and it is appropriate to summarize the totality of the software tasks presently forseen to give a sense of the magnitude of the required effort.

During the precommissioning of L1CalTrack and well in advance of the general installation shutdown, first using muon boards then production L1CalTrack boards, a substantial inventory of software must be prepared and installed:

a) Modify VxWorks to work with L1CalTrak
b) Add L1CalTrack startup to cold-start GUI's
c) Generate first pass trigger logic for L1CalTrack (e.g. BOT, L1CTT-only, L1CAL-only, L1CTT and L1CAL for electrons, taus, jets).
d) Modify COOR download to include L1CalTrack
e) Add L1CalTrack to parity-check GUI
f) Add L2CalTrack to mtm-term GUI
g) Add monitoring histograms to L3 Examine
h) Add L1CalTrack to online simulator-hardware comparison code
i) Add L1CalTrack to online efficiency code
j) Develop code to readout L1CalTrack data from RDS
k) Develop code to write L1CalTrack results to TMB
l) Complete L1CalTrack simulator in tsim_l1muo add DFEB-CTOC chain to trigsim, both for doublet-based (Run2a) and singlet (Run2b) logic

In fact, task l) is the purview of the Upgrade Project and the effort and timelines for its management are overseen by the Project. The SC-IPC has tabulated the effort and timelines only for the tasks a) through k) in its analysis of the software effort required for RunIIb (see section 3.2).

As a result of the precommissioning work for L1CTT with preproduction/production L1CTT DFEB's, software tasks such as the following also will have been achieved prior to the general installation shutdown:

a) update DFE power supply control for DFEB crates
b) raw Ethernet driver (done)
c) raw Ethernet serializer
d) add DFEB (+extra CTOC) boards to dfe_ware database
e) add download interface for DFEB to dfe ware/EPICS (Fall 2004 tests can proceed without this)
f) update link test scripts for Mixer-DFEB-CTOC connection test
g) create DFEB firmware for Mixer-DFEB-CTOC link list
h) update trigsim, both for old and new equations
i) add DFEB + new CTOC to CTT_Examine
j) add DFEB + new CTOC to CTT offline verification programs

During the precommissioning phase those who have developed this software will install and debug it as they accomplish the following operational tasks:

a) establish communication to the new crate controller (CC) via SCL and GbE
b) establish CC operation
c) establish communication with DFEB
d) establish basic DFEB operation
e) load doublet equations into DFEB
f) load CTOC (firmware accommodates missing inputs)
g) load CTTT (optional)
h) load CTM [Central Track Manager, the official name for the Muon Trigger Manager of the CTT] with new CTTT map (terms 176—191) (optional)
i) run link tests for all new connections (check LVDS input timing with scope?)
j) For standard thresholds, AFE test vectors, and low thresholds: verify I/O term rates; compare to old system (other inputs disabled); verify 0x13 vs old system; verify 0x13 output against existing trigsim (CTT_examine); offline efficiency analysis;
k) Load singlet (=new) equations: verify 0x13 output vs trigsim (for this we need the new equations in trigsim); offline efficiency analysis
l) Add modules and repeat above testing; adjust firmware for rates vs. efficiency

The SC-IPC assumes the development of this software and the precommissioning of it and the new DFEB's is managed by the Upgrade Project and does not include it in the effort summaries tabulated in section 3.2.

4.3 Working Group 3: L2 βeta Upgrade, STT

The L2 βeta processors are scheduled to be installed "adiabatically" in the L2 trigger system as they become available in advance of the start of RunIIb. Splitters and fibers for STT are also scheduled to be installed in advance of the start of RunIIb, so that when the shutdown for general installation begins, the balance of the new STT hardware is installed in the first week of the shutdown and technically commissioned after the operation of the new L1CTT begins. Full physics commissioning of STT will be completed by week 21 after the beginning of the shutdown.

4.3.1 *Effort & Schedule of Installation and Commissioning*

The development of the L2 βeta and STT hardware and the firmware for the STC, software updates for the TFC to accommodate the Layer Zero silicon, and the software for downloading and controlling the systems, is managed by the Upgrade Project and these activities are not the purview of the SC-IPC. Care will be needed to ensure adequate effort is allocated by the Upgrade Project to the development of L2 algorithms.

A resource loaded schedule for the installation and commissioning of L2 βeta and STT exists and it forms the basis of the SC-IPC analysis presented herein. The results of this analysis is tabulated in the Executive Summary, section 3.3. The total effort required for installation and commissioning is small compared to the other upgrades.

4.4 Working Group 4: Trigger Simulation

4.4.1 *Effort & Schedule of Development and Installation of a Trigger List*

A detailed planning exercise is underway to provide a basis for estimate of the effort required to develop the RunIIb trigsim system, and the trigger list and prescale set required for RunIIb physics data-taking.

The period of development for Trigsim is conveniently divided into three segments: the integration of code for the upgraded detector and trigger; the testing and commissioning of the software; the development of the trigger list.

In the tables presented in this section the following abbreviations have been used:

- *exp* : expert in that particular subsystem. Should have significant knowledge of the subsystem and the structure of the simulation code
- *ne* : non-expert. Should have some experience in running D-Zero code. Does not need any detailed knowledge of the subsystem
- *ne+exp* : non-expert with significant supervision of a subsystem expert.
- *trig exp*: Expert with detailed understanding of the trigger system but need not have any particular software skills of knowledge of the trigger simulator
- *tb + pg* : trigger board and physics groups.

It is assumed in general that *gs* and *exp* will have average or better C++ coding skills.

4.4.2 *Integration of Code for the Upgrade*

In Table 7 is indicated the timeline and effort schedule for the integration of code for the upgraded detector and trigger.

Trigsim Development				
Task	**Begin Date**	**Must Finish By**	**Resource**	**Level**
L1CalTrig Integration	underway	1-Jan-05	Upgrade Project	exp
L1CTT Integration	underway	1-Jan-05	Upgrade Project	exp
L1CalTrack Integration	underway	1-Jan-05	Upgrade Project	exp
Offline L1 Unpacking	ASAP	1-Jan-05	1 MM	exp
L1 => L2 Integration	ASAP	1-Jan-05	1 MM	exp
Layer 0: GEANT	ASAP	1-Jan-05	2 MM	exp
Layer 0: Unpacking	ASAP	1-Jan-05	1 MM	exp
Layer 0: Calibration Code	ASAP	1-Jan-05	1 MM	exp
Layer 0: L3 Unpacking	ASAP	1-Jan-05	0.5 MM	exp
Layer 0: L3 Tracking	ASAP	1-Jan-05	4 MM	exp
L2 STT	ASAP	1-Jan-05	1 MM	exp
Preliminary trigger terms	underway	1-Jan-05	1 MM	trig exp

Table 7. Integration of Trigsim Code for the Upgrade

Note that the first three items listed in Table 7 are the purview of the Upgrade Project and the effort and timelines for their management are overseen by the Project. The SC-IPC has tabulated the effort and timelines only for the balance of the tasks in the table in its analysis of the software effort required for RunIIb.

In Table 7 almost every task requires the attention of an expert. It is possible in those cases listed as *exp* that some of the expert's contribution could be replaced by some non-expert post-doc or graduate student. However, this will probably result in some delay to the schedule. Since most of these tasks involve the writing of significant chunks of C++ code it is possible that the bulk of the coding could be carried out by the expert with the non-expert contributing mainly in the testing phase. How much of a delay this causes will depend entirely on what skills the non-expert brings to the task. Someone familiar with the D-Zero coding and runtime environment might reasonably take a few days to get up to speed. A completely green graduate student with little or no coding experience might require at least a month to be able to contribute in a useful manner.

The January 1st deadlines are driven by the requirement that the tools required for trigger list development be available and certified for use by March 1, 2005. However, this does not necessarily represent the end of all algorithm development for the various upgrade components but rather a checkpoint in the development. At this point all components should be integrated in some fully functional form but perhaps not implementing the absolute final design of the algorithm. This provides a baseline release of the code which can be used to begin work on the trigger list as well as a stable code base from which to continue with algorithm development. It is clear that the closer the code is at this point to the final algorithm the quicker the trigger list development can be expected to converge to a final trigger menu.

The schedule outlined above does not have much room for delay. The level 3 components could, in principle, be allowed to slip until June 05 and then be certified since they are not specifically needed by the trigger board until the initial L1/L2 trigger menu has been completed. It may also be possible for the trigger board to converge somewhat faster on this list than was the case for v13.

The most worrisome pieces of the overall trigsim schedule are in Table 7. It is important that those items marked with begin dates "ASAP" be addressed as soon as possible. Particularly critical are the modifications to level 2 to handle the new outputs of level 1. While there is a relatively low risk alternative in the case of level 3 – simply ignore the new layer 0 – it is clearly desirable to have tracking triggers which can take advantage of the upgrade. It is important to recognize that this work should begin as soon as possible.

4.4.3 *Software Testing and Commissioning*

The TrigSim commissioning period covers the certification of the code base and the completion and testing of all tools required by the physics groups and the trigger board for the development of the post-2005-shutdown physics trigger list. The timeline and task list are presented in the Table 8 below.

Trigsim Commmissioning (Test Release Frozen 03 Jan 05)				
Task	**Begin Date**	**Must Finish By**	**Resource**	**Level**
Basic bugfixing	3-Jan-05	17-Jan-05	1 MM/subsyst	exp
Check Outputs	17-Jan-05	1-Feb-05	1MM/subsyst	ne+exp
Default trigger list	3-Jan-05	1-Feb-05	0.5MM	exp
Milestone 1: Cut Production Release: 1-Feb-05				
Full Certification	1-Feb-05	1-Mar-05	1 MM/subsyst	exp
Ancillary Tools				
Trigsim cert	3-Jan-05	1-Jan-05	0.5 MM	ne+exp
Ratetool	1-Feb-05	1-Mar-05	1 MM	exp
Milestone 2: All Tools Certified and ready for use: 01-Mar-05				

Table 8. Trigsim Commissioning

The first two steps of the commissioning task list constitute pre-certification. Here basic debugging is carried out through code review and simple testing of the algorithm outputs. It is estimated that with 1 FTE/subsystem at the expert level this step can be completed in 4 weeks, allowing the production code release ("p18"?) to be cut from the test release on 1st February. An important step in preparing to carryout the full certification is the production of a default trigger list. This is a list which contains references to every algorithm at each level of the trigger with some default set of very loose trigger cuts. The purpose of this list is to ensure that all trigger terms are exercised and provide lists of trigger objects which can be examined in the output.

Once the release tags have been created full certification, involving detailed studies of the algorithm performance, can begin. It is anticipated that there will be several changes to the level 3 system over the coming year, involving algorithmic improvements to address the challenges posed by the predicted increases in luminosity. This implies that essentially the whole level 3 code base will have to go through the certification process. For level 1 and level 2 it is conceivable that only those subsystems which have changed and the components dependant upon them will require certification. In each case it is estimated that 1 FTE/subsystem will be required to achieve the second milestone by 1st March 2005.

Two ancillary tools for processing the output of the trigger simulator: *trigsimcert*, and the *ratetool* are useful in the context of the certification and trigger list development respectively. The first of these provides a root-based format containing all the outputs of the trigger simulator for analysis; the second allows the construction of sample trigger terms and estimates of their rates based on the contents of the L1L2Chunk and L3Chunk in the DST or TMB output formats. In order for certification to proceed in a timely manner it is important that *trigsimcert* be upgraded to match the corresponding changes in the trigger simulator. The *ratetool* has proved essential in the development of the v13 trigger list and it is anticipated that this will continue to be the case in the construction of future global physics trigger lists. It is required that this tool be updated to allow the prediction of trigger rates for the upgraded trigger system. The SC-IPC believes this work for L1Cal is nearly complete and it must be completed no later than 1st March 2005 in order to meet the TrigSim commissioning milestone. The need for L1Cal to be available in *ratetool* by that date is undisputed in order to have a physics triggerlist by the end of the installation shutdown. The SC-IPC believes L1CTT and L1CalTrack will also be incorporated by early 2005, but if L1CalTrack were delayed, then it might be possible to introduce it in a later version after data-taking is underway.

4.4.4 *Trigger List Development*

This section describes the work needed to build and test a full trigger list for collecting physics quality data. The timeline (assuming week 1 of the shutdown is the last week of June), 2005) and task list are show in Table 9 below.

Trigger List Development				
Task	**Begin Date**	**Must Finish By**	**Resource**	**Level**
L1/L2	1-Mar-05	Week 6	5 MM	tb+pg
L3	1-Mar-05	Week 11	6 MM	tb+pg
DB Entry + checking	Week 11	Week 15	1 MM	trigmeister
Download testing	Week 15	Week 16	1 MW	exp
Post-shutdown testing	First data	Week 18	2 MM	exp
Milestone 2: First run taken with RunIIb trigger list two weeks after first data				

Table 9. Trigger List Development

Work on the v13 trigger list is essentially complete and work has already begun to prepare the next generation of trigger lists. However, the real work of studying the trigger rates and overlaps for the upgraded system will not begin until the trigger simulation and *ratetool* have been certified. Based on previous experience it is estimated that from this point between 4 and 5 months will be required to finalize the initial L1/L2 trigger menu. The level 3 trigger list will then follow after about 1 month's additional work. A further 6 weeks are then needed to actually enter the trigger list into the database and perform all the appropriate cross-checks and download tests to make sure all triggers have been included correctly.

Trigsim validation – checking that hardware readout agrees with simulation/emulation -- begins in steps as each new element of the upgrade is operational, i.e. when the end of the technical commissioning of each is completed, as detailed in 3.0. These dates are contained within the "Post-shutdown Testing" period shown in Table 9.

Before the full commissioned trigger list can declared operational it will be necessary to collect a few hours of enhanced bias data with Layer Zero in the readout to provide a sample for refining the silicon tracking based components of the trigger list. It is estimated that 2 weeks will be required to analyze the data and feedback the results into the trigger list and prescale sets. This last does not delay taking physics-quality data; real physics menus that read out Layer Zero but do not use it at L2 or L3 for some time will not cost luminosity.

The SC-IPC trigsim working group firmly believes that in order to complete this work on schedule it is vitally important that the correct experts be identified and assigned to work on the appropriate tasks. Non-experts working on these tasks need to have an appropriate level of supervision and oversight and must possess sufficient computing skills to be able to contribute effectively.

The trigger simulator remains an expert tool. There has been some discussion within the SC-IPC as to whether this should continue to be the case or if there is work that could be done to improve the ease-of-use for this program. This only falls within the purview of the SC-IPC insofar as it may affect the length of time taken for the trigger board to converge on a trigger list.

If the simulator is easier to use then it might be expected that relevant results could be produced faster, with fewer errors and with less expert intervention. It is also clear that a piece of software with more widespread use will generally have its programming mistakes revealed sooner than otherwise.

A resource loaded schedule for the development the RunIIb trigsim and the trigger list and prescale set needed to resume physics operation after the installation of the upgrades exists and it forms the basis of the SC-IPC analysis presented herein. The results of this analysis is tabulated in the Executive Summary, section 3.4.

4.5 Working Group 5: Layer Zero Silicon

4.5.1 *Effort, Cost and Schedule of Installation and Commissioning*

Substantial planning has been conducted to provide a basis for the estimate of the effort required to install and commission the new Layer Zero system, and to develop the timeline that will govern the execution of these tasks. The schedule resulting from this planning is reasonably mature, particularly with regard to the hardware installation tasks.

The schedule can be grouped into three main sections: pre-installation, installation, and commissioning. The pre-installation activities consist of modifications to online silicon software, modifications of offline silicon software, and miscellaneous hardware infrastructure tasks to be completed during the Fall 2004 shutdown. Technical commissioning consists of demonstrating the full operation of all channels from the control room prior to closing the detector, after the system has been cabled and cooled to operating temperatures. All online software required to operate the system must be at hand for this task, which will prove its proper function.

Layer Zero physics commissioning begins after the return of the Tevatron beam. This commissioning is broken up into two sections. The first is hardware physics commissioning, involving timing and noise studies. Data taken during this period of four weeks will not be physics quality. At the conclusion of hardware physics commissioning, the Layer Zero detector will be ready to take physics quality data. The second period of Layer Zero commissioning is algorithms physics commissioning. All offline software must be available prior to this commissioning. The work involves study of the new software algorithms using physics quality data. The final item of algorithms commissioning will be the production of simple physics distributions that will confirm that the Layer Zero data and reconstruction are physics quality and suitable for analysis.

The Layer Zero Upgrade Project acknowledges the importance of developing the necessary online and offline software that will be needed to bring Layer Zero to full physics operation. Some software packages, e.g. unpacking and tracking utilities needed for trigger simulation, will be developed by the trigger simulation group. The modification of all the online code packages to include Layer Zero, and the offline code modifications to incorporate Layer Zero into D0GSTAR and RECO, must be completed well in advance of the 2005 shutdown to make installation at that time meaningful.

Among the online software packages that must be modified are the tools needed for the new SVX4 devices in Layer Zero:

a) Download GUI
b) Calibration database
c) Calibration procedures
d) Online monitoring
e) High voltage GUI

Among the major groupings of offline software projects are:

a) Simulation
b) Unpacking and calibration
c) Cluster reconstruction
d) L3 trigger algorithms
e) Track reconstruction
f) Monitoring (examine, event display)

Significant portions of the software including simulation and unpacking must be ready by December 2004 when the prototype SVX4 HDIs will be included in the D0 data stream.

The specific work needed for trigger simulation occasioned by Layer Zero is referenced in the Trigsim summary section (7.4.1), and it includes

a) New geometry in D0GSTAR
b) New SMT hit storage interface
c) New SMT hit digitization package
d) New SMT hit unpacking to raw data chunks
e) New SMT cluster reconstruction packages
f) New RECO geometry and track reconstruction

It is estimated that 6 FTE physicists must be recruited immediately in order for the silicon online and offline software to be ready for Layer Zero next year. Perhaps half of these are required for trigsim-related tasks.

A resource-loaded schedule for installation and commissioning of Layer Zero silicon exists and it forms the basis of the SC-IPC analysis presented herein. The results of this analysis are tabulated in the Executive Summary, section 3.5.

4.6 Working Group 6: AFE II

4.6.1 *Effort, Cost and Schedule of Installation and Commissioning*

As noted in section 2.6, when the AFEII project completes a prototype system in late 2004/early 2005, the SC-IPC will prepare an analysis of the costs in effort and time required to install and commission the new system.

4.7 Working Group 7: Online

4.7.1 *Effort, Cost and Schedule of Installation and Commissioning*

As noted in section 2.7, all of the activities of the Online upgrade are managed by the Upgrade project and none require further comment by the SC-IPC.

5.0 References

1. D0News, General Folder, April 20, 2004
2. See for example, “The RunIIb D0 Detector Project Management Plan”, October 28, 2002
3. D0 Layer Zero Conceptual Design Report, October 2003, D0 Note 4415; online at http://d0server1.fnal.gov/projects/run2b/silicon/www/smt2b/Layer00/l0.htm
4. “Algorithms and Architecture for the L1 Calorimeter Trigger at D0 Run IIb”, J. Bystricky, et al., *IEEE Transactions on Nuclear Science*, Vol. 51, No. 3, pp 351—355, June 2004.
5. “The D0 Central Track Trigger”, J. Olsen, et al., *IEEE Transactions on Nuclear Science*, Vol. 51, No. 3, pp 345---350, June 2004.
6. “The Run IIb Trigger Upgrade for the D0 Experiment”, M. Abolins, et al., *IEEE Transactions on Nuclear Science*, Vol. 51, No. 3, pp 340—344, June 2004.
7. “Run IIb Upgrade Technical Design Report”, Fermilab Pub-02-327-E, Sept. 2002.

6.0 Charge to the Standing Committee

Folder: GENERAL
From: blazey@dsc09-waf-dc-206-214-1-169.rasserver.net
Subject: Standing Committee on Upgrade Installation-to-Physics Commissioning
Date: 20-APR-2004 22:05
Expires: 04-JUL-2004 22:05

Dear Colleagues

We are pleased to announce formation of a Standing Committee on Upgrade Installation-to-Physics Commissioning. The committee is being charged with developing an integrated experiment-wide plan for the installation and full commissioning of the upgrade elements. Their findings will be used as input to the spokesperson's deliberations as we decide on the scheduling of, and the collaboration-wide resources required for, the installation and commissioning of each of the upgrade components. The full charge to the committee is attached below.

The current committee membership represents the unusually extensive nature of the task. It consists of Ela Barberis, Volker Buescher, Bill Cooper, Stefan Gruenendahl, Bob Hirosky, Rick Jesik, Ken Johns, Eric Kajfasz, Breese Quinn, Lisa Shabalina, Gordon Watts, and Taka Yasuda, and will be co-chaired by Cecilia Gerber and Rich Smith. A small number of members may be added as needed, to bring in additional expertise. The Committee will be necessarily be drawing on the expertise of many in the collaboration; we ask each of you to assist them when called upon. We thank the committee members for agreeing to take on this important and challenging task.

Sincerely

Jerry Blazey and John Womersley

****

Charge to the Standing Committee on Upgrade Installation-to-Physics Commissioning

In order to continue to exploit and maximize the physics potential of the experiment throughout Run II, the upgrade trigger, DAQ/online, and silicon sub-projects must be installed and commissioned in a manner that to the greatest extent possible minimizes the total downtime of the experiment. The installation and technical commissioning of the hardware represents a portion of that process; full commissioning of the elements further downstream (physics commissioning) will also be required in order to resume the acquisition of physics-quality data. A collaborative effort that effectively bridges the usual hardware/software boundaries is essential in order to optimally coordinate this process. To that end, we have assembled a Standing Committee that is being asked to generate a detailed plan describing the evolution of the various activities. This plan will include, but not necessarily be limited to, the following areas:

- Installation and technical commissioning of the hardware;
- Calibration databases, both online and offline;
- Updates to data unpacking & formatting;
- Development of Level 3 filtering algorithms;
- Updates to clustering and track finding, both online and offline;
- Development of Monte Carlo and TRIGSIM.

We have asked the Committee to come up with a plan that includes a breakdown and timelines for each of the major efforts, the resources required (including manpower type – physicist, engineer, technician, etc.), and the sequence in which the various pieces need to be developed. Given the complexity of this set of processes, we have assembled a broad group of experts to develop this plan. The Committee therefore includes personnel from the upgrade project, as well as those more closely affiliated with areas beyond the upgrade and hardware efforts.

We imagine that the work associated with this effort will go on until the upgrade elements are commissioned, but that the studies will be done in more focused periods, not as a constant effort. It is understood that the Committee will work closely with the Algorithms Coordinator, the Technical Coordinator, and the Run IIb Project Manager, taking direction and guidance from them in developing the plan. The spokespersons, in conjunction with the aforementioned managers of the experiment, will be asking others in the collaboration to assist the Committee with any studies that are needed, and to generally lend their input and expertise as the plan is developed. We have also asked the co-chairs to provide periodic oral reports on their progress in various forums, including collaboration meetings and the All DZero Meeting. A written report describing the findings of the Committee will be updated and submitted to the spokespersons approximately every three to four months, and will be used to evaluate the optimal approach to, and timing for, the installation of the upgrade elements. We have requested a first update be presented at the June 2004 workshop in Fresno.

This is a most challenging task, and we very much appreciate that this group has agreed to serve the experiment in this capacity. We have every hope and expectation that the timely establishment of this Committee, and the more integrated approach to installation and commissioning that their efforts represent, will enable us to most fully optimize the physics reach of the experiment during the coming run.

Jerry Blazey and John Womersley

7.0 Description of the Upgrades

The separate elements of the Upgrade Project are described in detail in available documents for the Layer Zero silicon [3], the L1CalTrig [4], the L1CTT [5], and the L2Beta and L2STT upgrades [6,7]. In what follows, brief summaries of the purposes of the upgrade elements are outlined to provide a setting for the work of the SC-IPC.

7.1 Summary of the L1CalTrig Upgrade System (Working Group 1)

It is not feasible to accommodate the higher luminosities anticipated in Run IIb by increasing trigger rates. The L1 trigger rate is limited to ~3--5 KHz for acceptable deadtime (L1 hardware tests take ~ 5 microseconds, readout after a L1 trigger ~ 10 microseconds). Approximately 80% of the L1 rate comes from calorimeter-based triggers. To obtain increased rejection of QCD jet backgrounds and add new tools for the recognition of Higgs, supersymmetry, or other new physics at large mass scale, a family of upgrades to the L1CalTrig is being prepared to

- Sharpen the turn-on for jet triggers, reducing the rates
- Improve trigger turn-on for electromagnetic objects
- Add shape and isolation cuts to electromagnetic triggers, reducing the rates
- Add the ability to match tracks to energy deposition in calorimeter trigger towers, reducing the rates
- Add the ability to include the ICR (Inter Cryostat Regions) energies to the calculation of jet energies and missing ET
- Add the ability to form topological triggers to aid on triggering on specific final states

The RunIIb upgrades which replace the existing calorimeter trigger to facilitate these improvements consists of:

- Digital filtering with matched filter and peak detector algorithms to allow proper triggering on correct bunch crossings
- Sliding window trigger tower algorithms to sharpen identification of physical objects
- The addition of ICR layers to the global ET sums

The hardware to support these new facilities consists of:

- ADC-Digital-Filter (ADF) boards that receive the analogue trigger tower signals from the existing calorimeter BLS cards, digitizes and converts from energy to ET, and exercises digital filtering to associate with the correct bunch crossing.

- ADF Timing Fanout boards that send trigger framework timing signals to the ADF cards

- Trigger Algorithm Boards (TAB) that receive the trigger tower transverse energies from the ADF, produce the EM and jet cluster ET's using the sliding window algorithm, and begin the global summing for scalar ET total and MPt.

- Global Algorithm Board (GAB) that receives data from the TAB's and produces final ET total and MPt, plus interfaces to the existing trigger framework and a timing fanout.

7.1.1 *Design Features of L1CalTrig which Mitigate Risk*

No ASIC's are required for the L1Cal trigger electronics. Sufficient spares will be available to promptly replace problem cards. A Test Waveform Generator (TWG) provides an alternate source of input data for testing.

The system interfaces to the D0 slow controls system so that software triggers can be generated for testing.

To reduce cabling vulnerabilities, the ADF cards are 6U VME (32 channels) with all analogue and digital I/O connectors on the rear side of the card. Analog inputs are via a J2 connector, with digital outputs assigned to the J0 connector.

7.1.2 *Status of Pre-Installation Risk Mitigation Planning and Activities*

In attempting to avoid a lengthy (> one year) commissioning period after installation, the L1CalTrig group will begin integration and precommissioning activities before installation of the L1Cal trigger in MCH1. A test site adjacent to the MCH on the sidewalk has been prepared with detector-based grounding, rack infrastructure, access to the SCL trigger/timing hub, and output from the Calorimeter BLS system. A eight splitter cards (four Trigger Towers per card) were prepared which duplicate the analogue signals of the associated BLS's; these were installed in the existing calorimeter trigger front end, and they bring actual D0 Run IIa calorimeter trigger tower data to the test site in real time. (N.B: the present limit -- 16 EM + 16 H -- to the number of split signals means the sliding window algorithms cannot be tested at the test site).

A special ADF timing/test card (SCLD) was prepared to link SCL timing to the ADF prototype and this has been tested at the test site. The VME/SCL card that distributes timing signals to and provides VME communication with the TAB's and GAB has also been tested at the test site. A TAB/GAB test card has been fabricated which was used to test ADF-TAB transmission before integration at the test site, and this same board was also used to test TAB-GAB transmission before the GAB prototype was sent out for fabrication. Data transfer from the ADF prototype to the TAB prototype has been achieved at the test site, as has output from a TAB to the L1Muon (Cal-Track trigger) and to L3. Quantities of data from the TAB have been sent to tape for analysis, and in the future the full chain BLS => ADF => TAB => GAB will be tested at the test site. These prototypes will be replaced with preproduction or production versions of the electronics in the fall of 2004.

As the production versions of the components become available, the full L1CalTrig system hardware (in five active racks, plus additional racks to accommodate the BLS-ADF patch panels) will be integrated and tested at the test site prior to Run IIb. The integration and testing must interface into the running experiment and permit numerous tests to be carried out without perturbing data-taking. The L1CalTrig team is planning to document to the collaboration that operation of the new L1CalTrig system will not perturb the system, and that data taken with it will not be worse than that from the present system. (This last goal is complicated by the fact that the limited number of splitter signals does not allow the sliding windows algorithm to be fully tested). Regarding the first point, operations at the test site will provide information about crashes/deadtime, downloading reliability, and show that monitoring tools, parameter determination, and unpacking for RECO are all fully operational. With the aid of data from the test site, plus Monte Carlo, rates & efficiencies can be predicted/verified, and L1, L2, and L3 trigger definitions (filter coefficients, thresholds, and/or terms, trigger list) be put in place and made operational.

When the RunIIb shutdown for installation begins the fully operational racks of equipment can be moved en masse into MCH1 during the Run IIa – IIb shutdown with minimal elapsed time.

7.2 Summary of the RunIIb L1CalTrack and CTT (Working Group 2)

The L1CalTrack Trigger will exploit matches in Φ of tracks from the L1CTT trigger with EM and jet objects from the L1Cal Trigger (see section 7.1 for a summary of the L1CalTrig upgrade) to reduce L1 trigger rates from EM and track triggers. Its inputs are therefore L1CAL, L1CTT, and FPS. It outputs information to the Trigger Framework, and data to L2 and L3 trigger system via the Muon L2 and L3 readout crates. The L1CalTrack hardware is very similar to the current L1MUON hardware, and the input connections to L1CAL and L1CTT are very similar, in fact identical copies in the CTT case, of the Run2a L1CTT – L1MUON connections.

The L1 Central Tracking Trigger (L1CTT) is part of the overall CTT system, and combines tracking information from the CFT and information from the central preshower system (CPS) to provide stand-alone track triggers for the identification of electrons and photons, plus track lists for other systems to perform track matching (e.g. L1Muon, L1CAL, and the L2STT). The individual portions of the L1 track trigger (CFT, CPS, FPS) are sent to L2 for trigger decisions there. In RunIIa track information from CFT consists of "doublet hits", i.e. an OR of signals from axial fibers in adjacent inner and outer fibers of a doublet layer, with a veto on one adjacent fiber to prevent the formation of more than one doublet hit per input singlet hit. For RunIIb the effective granularity of the CTT is increased by using information from individual CFT fibers rather than doublets. For high pt tracks all 16 layers will be used; for low pt tracks, high granularity is used in the inner 4 doublet layers and doublets in the outer low-occupancy layers. This hybrid scheme provides the best possible rejection for fake high pt tracks for running a high pt isolated track trigger, and yet it retains the ability to process low pt tracks for b-tagging in the STT.

To implement the new requirements, powerful new commercial FGPA's will replace those in use, and thus new daughter boards (DFEA, A for axial) on which they sit, are required, for the digital front end (DFE) system. In total two of the seven CTT crates will be replaced.

In what follows, the L1CalTrack will be treated separately from the L1CTT for textual clarity.

7.2.1 *Design Features of the L1CalTrack Trigger which Mitigate Risk*

The L1CalTrack Trigger uses existing L1 Muon architecture (in operation during Run IIa for ~3 years already) with small modifications. The L1 Muon trigger matches Φ of tracks from the L1CTT with muon objects, a function entirely similar to that required for a L1CalTrack trigger. The adaptation of existing designs for the Muon Trigger Crate Manager (MTCM), Muon Trigger Cards (MTCxx) and Muon Trigger Flavor Boards (MTFB – to be renamed Universal Flavor Boards -- UFB), means issues such as synchronization, buffering, outputs to L2 and L3, electronics testing, operational monitoring, power supplies and rack infrastructure, all have proven, working solutions.

7.2.2 *Status of Pre-Installation Risk Mitigation Planning and Activities: L1CalTrack Trigger*

Detailed planning is underway to integrate the L1CalTrack Trigger system into the experiment in a step-wise fashion. In mid-2004 effort will be expended to generate Beginning-Of-Turn (BOT) triggers from the trigger framework and to readout both L1CalTrack crates in MCH1.

Initially, existing spare MTCM's and MTCxx's will be installed, to be replaced later with production L1CalTrack cards as they become available (in late 2004). (All production MTCM, MTCxx, and MTFB's will have been tested at Arizona and Fermilab using the Muon Test Trigger (MTT -- this system, developed at Arizona and certified and already incorporated in loop tests there, permits the testing of all inputs, timing, buffering, UFB's, and L1/L2/L3 outputs. Soon five MTT's will be available, to be installed at Arizona, Boston, and Fermilab). The new MTCxx boards will conform to JTAG standard (IEEE 1149.1) so boundary scan testing is possible.

Initially, spare L1MU inputs (on cables FPD4 and 5 which run from PE to MCH1, one of which is connected to an A-layer PDT) will be used to generate the BOT triggers to send to the TF. Demonstrating a stable BOT rate of 47712 Hz is an excellent demonstration that the system is working. A second important test is to readout the L1CalTrack crates into L3. Both the bunch crossing numbers and trigger outcomes can be checked. This testing can be done between stores.

Next, the L1CalTrack hardware will be replaced with L1Muon hardware. This demonstrates that the L1CalTrack MTCM's and MTCxx's are functioning properly in the running experiment.

It is important to realize that further testing must await the arrival of the new L1CTT and L1Cal inputs. The latter do not become available until the RunIIb shutdown is well underway.

7.2.3 *Design Features of the L1CTT which Mitigate Risk*

As commissioning of the RunIIa L1CTT matured, it was realized that additional features for the RunIIb upgrade hardware for L1CTT were strongly indicated. New DFEA backplanes and crate controllers are seen as crucial for facilitating downloads, and new DFE motherboards plus LVDS splitters are seen as crucial for integrating the system into D0. The new motherboards boards will conform to JTAG standard (IEEE 1149.1) so that testing is facilitated, and the DFEA daughterboards have been folded into the motherboards. I/O seed and capture buffers will be designed into the new DFEA cards to facilitate testing.

The addition of the Gbit/s Ethernet interface to the crate controller is intended to facilitate downloads, which presently can take about four hours (with two parallel jobs). This scales to 12 hours for a download of the RunIIb equation files over the existing 1553 bus. Given the need for restarts to overcome uncorrectable download errors, a 1553 download might never finish successfully. This interface entails some risk so a parallel port will also be added that will be at least as fast as the 1553 bus, and can be used for bench tests and whenever there is access to the system.

During the fall 2004 shutdown the installation of a new crate and power supply on the platform is scheduled, with LVDS splitters and cables. The new crate controller and prototype DFEA's can be added (and replaced) during minor shutdowns so that component testing can occur. This will then allow to test the new system with real data, and with standard output connections, in parallel to operation of the CTT.

For the new DFEA crates the input cable plant will be moved from the front of the boards to the rear (transition boards will no longer be needed) and the new backplanes will be designed so the input and output cables are properly separated to facilitate debugging of cables. LED indicators will be provided on the front panels. A 48 V power supply will be used to reduce the currents required at the crate.

7.2.4 *Status of Pre-Installation Risk Mitigation Planning and Activities: L1CTT*

The addition of LVDS splitters is key to the planning for pre-installation risk mitigation for the CTT. During the 2004 fall shutdown the SCL will be routed from MCH1 to PC03 on the platform (and GbEthernet from FCH2 to the platform), a new crate and 48 V power supply and backplane installed on the platform, and the LVDS cable extensions and splitters installed. A secondary downstream chain consisting of one additional CTOC and one CTTT will be added to PW03-2, and the necessary cables installed. These boards, together with the prototype DFEA crate and a few VRB readout channels in crate x13, form a complete L1CTT chain for testing of the Run2b system. Later, short accesses can be made to install the new crate controller and DFEA prototypes ("DFEB") for testing.

Failure to demonstrate the operation of the splitter system in time jeopardizes this schedule and might lead to the delay of the installation of the new L1CTT system until the 2005 shutdown, and the commissioning of it after RunIIb begins, reminiscent of the way in which the RunIIa L1CTT was commissioned.

7.3 Summary of the RunIIb L2 βeta & L2 STT Upgrades (Working Group 3)

7.3.1 *Summary of the RunIIb L2 βeta*

At Level 2, preprocessors analyze output from the L1 trigger of each detector system in parallel, including calorimeter, preshower, fiber tracker, and muon components. The STT also examines hits on the SMT barrel sections. The input rate (3 kHz) to L2 is limited by the silicon digitization deadtime, and the output rate (1 kHz) is limited by the full muon readout and calorimeter precision readout deadtime. To accommodate the higher luminosity of RunIIb while maintaining the same rejection factor of 5, more powerful processors for the L2 system are desired, which can execute more complex algorithms.

Algorithm changes forseen for Run IIb, such as incorporating vertex information from STT tracking information (vs. assuming all Cal Tower information is relative to a fixed vertex at Z = 0) in order to enhance L2 resolution for jets, EM objects, and Missing ET, will allow higher trigger thresholds for a given efficiency. Vertex resolution of 10 cm can improve trigger rates by 30-60% depending on pseudorapidity constraints, providing a rejection of 1.4 – 2.5 independent of L1. Another strategy is to add tower-by-tower calibration corrections (or thresholds in the case of missing ET) to improve the resolution of calorimeter objects.

The global processor, which combines the results of all the individual detector processors, to provide the final L2 trigger selection, can also benefit from increased CPU performance. By breaking the software restriction of the one-to-one mapping of the L1 trigger bits to the L2 bits, more trigger-specific processing can be applied to individual L1 contributions in L2, as is presently done in L3.

7.3.2 *Design Features of the L2 βeta Upgrade which Mitigate Risk*

The upgraded L2 βeta processor cards use existing βeta motherboards. A single board computer (SBC) resides on a 6U compact PCI (cPCI) card mounted on a 6U to 9U motherboard that supports the remaining functionality (FPGA and VME interface, and D0-specific Magic Bus and trigger framework communication) required to integrate the processor with the trigger data flow. Lower-performance SBC's can be inserted for Administrator-type processor cards, higher-performance (more costly) SBC's can be inserted where physics algorithms run. βeta processor cards using commercially available SBC's (e.g. Pentium III 1000 MHz processors) have supplanted original Alpha 500 MHz processors in the Run IIa L2 system. Higher performance SBC's have reached the market and the upgrade path for this system consists of replacing heavily loaded processors with higher performance models.

7.3.3 *Status of L2 Beta Pre-Installation Risk Mitigation Planning and Activities*

The proposed L2 Beta processors were originally conceived to provide alternatives to the Alpha processors which had lower than expected production yields, and to provide a clear upgrade path for the future. As such, the L2 Beta program is already well underway, and the new cards are scheduled to be installed one-by-one on an "adiabatic" basis well before Run IIb begins.

No new firmware updates for L2 Beta are required or planned. A scheme to evaluate L2 bit expansion for the online system has been adopted and effort for this has been added to the summary prepared by the SC-IPC. Effort for the development of new RunIIb algorithms is included in the Upgrade Project and is not recounted here.

7.3.4 *Summary of the RunIIb STT*

Raw data from the SMT from every L1 trigger accept is passed to the STT. Output from the L1CTT (tracks from hits on axial fibers of the CFT) is likewise sent to the STT which matches the L1CTT tracks with a subset of the hits from the SMT. Output from the STT consists of trajectories which greatly improve the precision of the measurements of transverse momentum and impact parameter made at L1. The STT also greatly reduces the fake rate of L1 tracks for which there are no SMT hits. The upgrades to the STT for Run IIb consist (except for the hotlink repeaters) of increasing the number of components in the existing system, something that is not expected to be difficult despite the increasing obsolescence of the existing components.

Including spares, twenty new VME motherboards are required, 70 STC cards (silicon trigger cards, which receive the data from the SMT front ends, filters it to associate hits with L1 tracks), 10 VTM (VME Transition modules, which split the optical fiber signals from the SMT to parallel paths to the SVX DAQ and another to the STT), 52 link transmitter boards (which have three LVDS serial link transmitters each for intermodule data transfer), 46 link receiver boards for intermodule data transfer, 20 BC (buffer controller) daughterboards which store data for L3 readout until a L2 trigger decision is made, 8 TFC (track fit trajectory) boards, 15 hotlink repeaters (the only component that requires design) which sends the data to the L2CTT, 13 optical splitters to create the data path to the STT, 52 optical fibers from the splitters to the VRB's and STT, and 125 cables for the LVDS serial links.

7.3.5 *Design Features of the STT Upgrade which Mitigate Risk*

As noted, the large number of components required for the STT are all essentially copies of the same components developed for Run IIa, except for the hotlink repeaters. The hotlink repeater boards are required to merge the outputs of two TFCs into a single hotlink output. They fit into the PC-MIP slots of the motherboard and have a PCI target interface to receive the output of one TFC and a hotlink receiver to receive the output of another TFC, some logic in an Altera FPGA to merge the two streams, and a hotlink transmitter for output of the merged data to L2CTT. All components are commercially available and the design is considered straightforward (only the transmitter is a new feature).

7.3.6 *Status of STT Pre-Installation Risk Mitigation Planning and Activities*

The readout of the new CTT frond end (DFEB's) into CTOC's will be verified prior to the 2005 shutdown. It has been suggested that a portion of the output from the new L1CTT can be included in the production readout during the 2004-2005 running period. Readout into STSX's will also be verified prior to the 2005 shutdown. Timing and data formats into the CTOC's and STOV/STSX's are identical, thus the competence of the new STT might also be made substantially credible in advance of the 2005 shutdown.

7.4 Summary of Trigsim (Working Group 4)

Trigger simulation for RunIIa has been developed essentially as an expert tool to simulate the D-Zero trigger system: hardware/firmware emulation for level 1 and software simulation at levels 2 and 3. Since the level 2 and level 3 triggers are software based essentially the same code can be used both online for triggering and in the simulator. For the most part the simulator has been used by experts to debug and evaluate the triggering performance of the various algorithms. Fundamentally, its inputs are data files containing a *Raw Data Chunk* and a trigger list, though in practice there are many other inputs which provide calibration constants, zero suppression thresholds, correction factors etc. The simulator includes in its output the triggering decision at each level of the trigger system and all the trigger objects used to make this decision. At the DST and TMB level this is contained in the L1L2Chunk and the L3Chunk. In the sections to follow the word subsystem will be used generally to refer to those separate components of each trigger level which reconstruct specific particle like objects such as Electrons, Muons, Jets, etc. , for example in level 3 there are subsystems such as : L3Jet, L3Electron, L3Tau, L3Muon, L3Track and so on.

As an expert tool the trigger simulator has generally required rather detailed prior knowledge – or a helpful expert on hand – to get useful information in the context of physics analysis trigger studies. Particularly problematic are the construction of useful trigger lists, ensuring that the correct configuration data for the task at hand is used and getting the outputs into a convenient format.

7.4.1 *Status of Trigsim Pre-Installation Planning and Activities*

The work to include the emulation of the upgraded level 1 trigger in the trigger simulator is already underway. L1CalTrack depends on the outputs of L1Cal and L1CTT so it cannot be fully implemented until these are completed. However, since the output format of L1CTT is not expected to change it should be possible to use the existing L1CTT emulation for initial testing of L1CalTrack. These tasks do not include the offline unpacking code required to extract the algorithm outputs from the raw data to be used in offline reconstruction and analysis. It is expected that the offline unpacking code (CVS package: *l1l2reco*) will require about 1 month of expert level work to write and test to production quality.

The output format of the L1Cal subsystem will change in the upgrade and so corresponding changes must be made in the level 2 trigger system. The same modifications are required in the online level 2 code and the trigger simulator (the code base is essentially the same).

The GEANT simulation of the new layer 0 silicon detector and the code for unpacking were enumerated in the trigsim section (4.4). The silicon unpacking, clustering and tracking code will need updating in the level 3 system to take into account the extra layer of detector. Preferably they would be updated to include the layer 0 in the tracking algorithm to improve the performance. However, if this is not done these components will still require modification to make sure the extra layer is quietly ignored. Estimates of the manpower required to do this range from 1 week to several months of expert time for ignoring the layer 0 or including it respectively.

7.5 Summary of Layer Zero Silicon (Working Group 5)

Given the decision by the directorate to cancel the construction of a full replacement for the RunIIa silicon tracker, the D0 team successfully defended the next best option – the addition of a “Layer Zero” detector in the existing radial gap (15 to 22.9 mm) between the new beryllium beam pipe already procured for RunIIb, and the support structure for Layer 1 of the existing RunIIa silicon detector. The resulting Layer Zero tracker provides a tracking layer at small radius with six phi segments and eight Z segments (48 x 256 channel HDI’s utilizing new SVX4 chips) with 98.4% phi acceptance.

The measurements of depletion voltage variation with dosage for the sensors in the existing detector indicate that the layer 1 sensors will begin lose significant numbers of channels at a total integrated luminosity of ~ 3-4 fb-1. Furthermore, at present approximately ~16% of the HDI’s on the detector are disabled. Studies show the Layer Zero addition will enable the recovery to non-irradiated values of the original b-tagging efficiency (or a per-jet increase by 15% in b-tagging efficiency for no irradiation). In addition, the improvement in vertex resolution afforded by the Layer Zero addition allows significant increase in the reach of B_s mixing: for reasonable predictions of a (nonzero) value of ΔM_s, the integrated luminosity required for a 3σ observation is substantially lower (e.g. at $\Delta M_s = 16ps^{-1}$, the luminosity required is reduced from approximately 3 fb^{-1} to 0.5 fb^{-1}).

7.5.1 *Design Features which Mitigate Risk*

At the time of the cancellation of the full replacement detector, the design of that device had matured to the point that sensors were in procurement and fabrication of ladders was imminent. The Layer Zero detector takes full advantage of this design progress (especially that for layers 1 and 2 of the original upgrade device) for sensor design, cabling, grounding, cooling, and support. As with the inner layers of the RunIIb design, the Layer Zero detector uses analogue cables to move the HDI’s to the ends of the detector to reduce the thickness in the central region.

In the light of known negative experience regarding noise in the CDF L00 silicon which also uses analogue cables, D0 has determined that poorly grounded conductors near sensors and cables are large sources of noise. The Layer 0 support tube is a carbon-fiber structure which has excellent electrical conductivity (comparable to copper at relevant frequencies). It was determined that forcing the carbon fiber support structures to ground potential via low-inductance connections is essential for providing a low noise environment. A solution was found by covering all carbon fiber surfaces with a laminated copper mesh giving noise isolation equivalent to a fully shielding Faraday cage.

A new feature not present in the RunIIb design is the pitch adapter, added to provide a wirebond platform between the sensor and the analogue cables. This platform eliminates the need for a separate bias voltage filter card and allows connections to the backs of the sensors to be moved to the ends thus reducing the effective width of the sensors and providing greater installation clearance. The pitch adapter also minimizes the number of different types of analogue cables required.

7.5.2 *Status of Pre-Installation Risk Mitigation Planning and Activities*

There are a number of activities that have been planned for the fourteen week shutdown beginning late August 2004 that help minimize the risk of installation of Layer Zero. Perhaps the most important is conducting a survey of the beam region apertures for Layer Zero installation. The device must fit into an extremely tight space and there are some uncertainties regarding the existing clearances, particularly near the Z=0 region. Verification of the Layer Zero design's compatibility with actual clearances is critical for ensuring a successful installation. A handful of prototype HDIs and the associated infrastructure will be installed within the detector volume during the Fall 2004 shutdown. These devices will be incorporated into the current online readout, enabling development and debugging of software modifications for the SVX4 chips. A proposal was made to remove the four H disks, refurbish the two inner ones, and replace them all during the Fall shutdown. The two outer H disks will be permanently removed when Layer 0 is installed, and there will be little time available to refurbish the inner disks at that time. However, constraints from other Fall shutdown activities such as the platform lifting, as well as the relative good shape of the inner disks relative to the outer disks makes the H disk refurbishment unlikely during the Fall shutdown.

Other work that will be done prior to the Summer 2005 shutdown includes expansion of the SMT bias voltage system, and installation of the new LVPS system for Layer Zero. Many new high voltage power supplies as well as new distribution boxes and cabling will be needed to provide bias voltage for the Layer Zero devices. This work can be done in approximately one week, mostly within the Movable Counting House, and therefore can be scheduled almost anytime between now and Summer 2005.

8.0 Acknowledgment

The SC-IPC is pleased to acknowledge the assistance of many in the collaboration who provided information and advice as this report was being prepared. Especially the SC-IPC benefited from the careful reading given this report by the Technical Integration Coordinators and many of those in the management structure of the Upgrade Project.